



US006148301A

**United States Patent [19]**  
**Rosenthal**

[11] Patent Number: **6,148,301**  
[45] Date of Patent: **Nov. 14, 2000**

**[54] INFORMATION DISTRIBUTION SYSTEM**[75] Inventor: **Joseph S. Rosenthal**, Atlanta, Ga.[73] Assignee: **First Data Corporation**, Atlanta, Ga.[21] Appl. No.: **09/108,796**[22] Filed: **Jul. 2, 1998**[51] Int. Cl.<sup>7</sup> ..... **G06F 17/30**[52] U.S. Cl. ..... **707/10; 707/3**[58] Field of Search ..... **707/3, 10; 345/333,  
345/327; 455/4.2****[56] References Cited****U.S. PATENT DOCUMENTS**

4,554,418	11/1985	Toy	179/2 DP
4,812,843	3/1989	Champion et al.	340/905
4,941,170	7/1990	Herbst	379/100.07
4,974,254	11/1990	Perine et al.	379/100
4,994,926	2/1991	Gordon et al.	358/400
5,131,020	7/1992	Liebesny et al.	379/59
5,291,302	3/1994	Gordon et al.	358/400
5,404,231	4/1995	Bloomfield	358/400
5,459,584	10/1995	Gordon et al.	358/434
5,537,586	7/1996	Amram et al.	707/3
5,555,100	9/1996	Bloomfield et al.	358/402
5,559,611	9/1996	Bloomfield et al.	358/407
5,649,186	7/1997	Ferguson	707/10
5,657,461	8/1997	Harkins et al.	345/333
5,754,938	5/1998	Herz et al.	455/4.2
5,754,939	5/1998	Herz et al.	455/4.2
5,835,087	11/1998	Herz et al.	345/327

**OTHER PUBLICATIONS**

*PR Newswire*, Aug. 11, 1982, (Distribution: To Financial).  
*PC Week*, Sep. 1, 1987, "RCA International service could cut Fax cost to Japan . . .".

*PR Newswire*, Oct. 13, 1987, "New System Cuts Costs, Increases Value of Fax Communications".

*Network World*, Oct. 19, 1987, "RCA Introduces Links, Fax Service and Gateway".

*InfoWorld*, Oct. 26, 1987, "Packer-Switched Network Updates Fax Technology".

*Data Communications*, Nov., 1987, "Fax Muscle".

*Network World*, Nov. 2, 1987, "U.S. Fax to Offer Packet Fax Service".

*PC Week*, Nov. 17, 1987, "Zap Mail's Loss is Fax Services Gain"; Section 2: Connectivity.

*Link-UP*, Nov., 1988, "Xpedite Launches Fax Software".

*Los Angeles Times*, Nov. 4, 1988, MCI Pioneering Separate Network for Fax Messages.

*PC Computing*, Mar. 1989, "The Fax is in the Mail".

Primary Examiner—Thomas G. Black

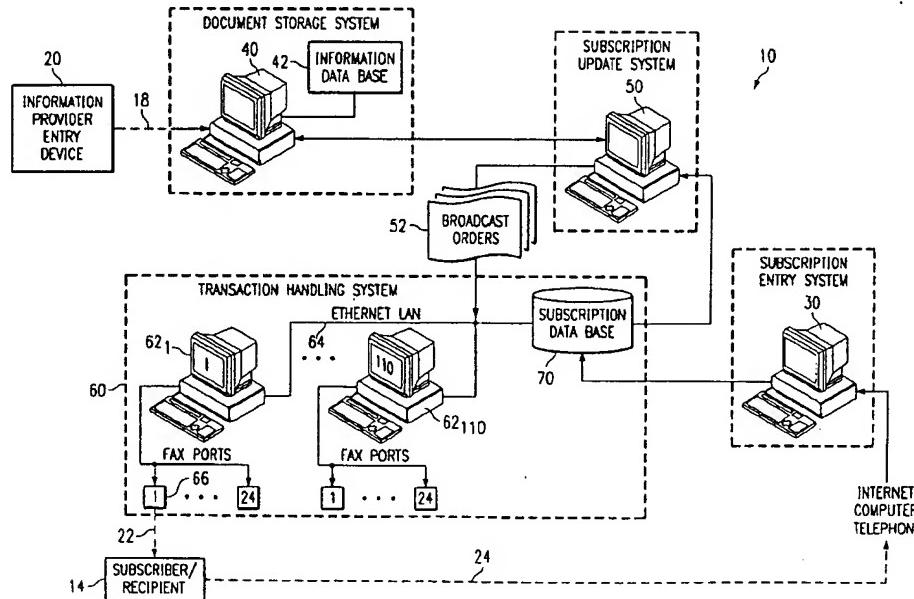
Assistant Examiner—Frantz Coby

Attorney, Agent, or Firm—Sidley & Austin

**[57] ABSTRACT**

A subscription facsimile system of the type in which providers supply new documents for storing the images thereof in a centralized database. Subscribers register with a system as to one or more subject matter areas of interest. A processor cyclically scans the document database for new documents that have been added thereto. The processor correlates the subject matter of the newly added documents to the identity of the subscribers having an interest in the same. The new documents are then transmitted by way of facsimile to the subscribers who have registered an interest in such subject matter.

**20 Claims, 4 Drawing Sheets**



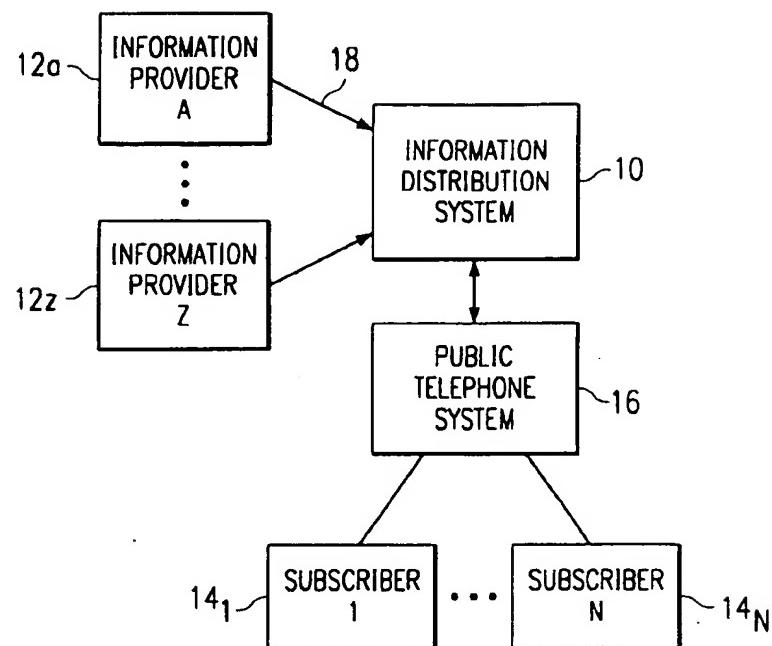


FIG. 1

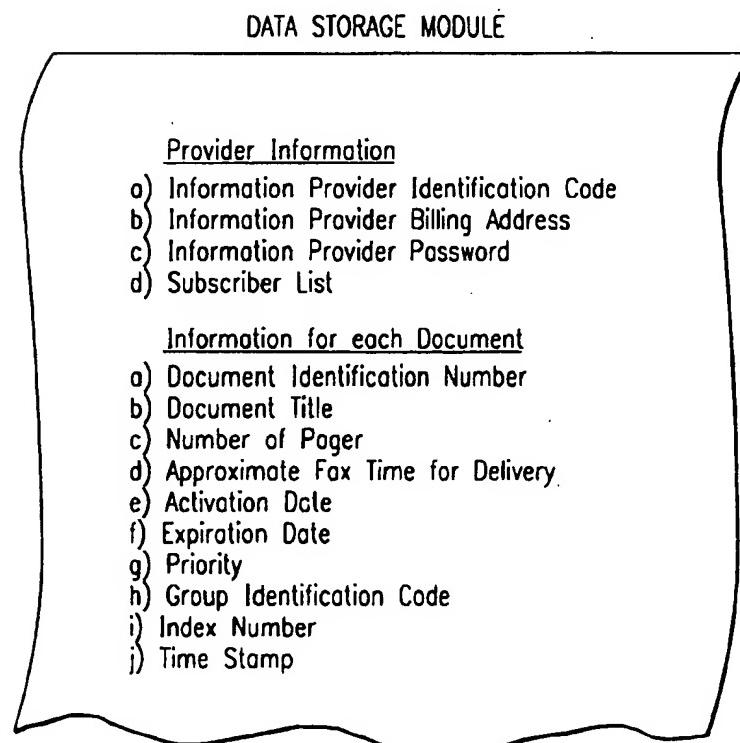
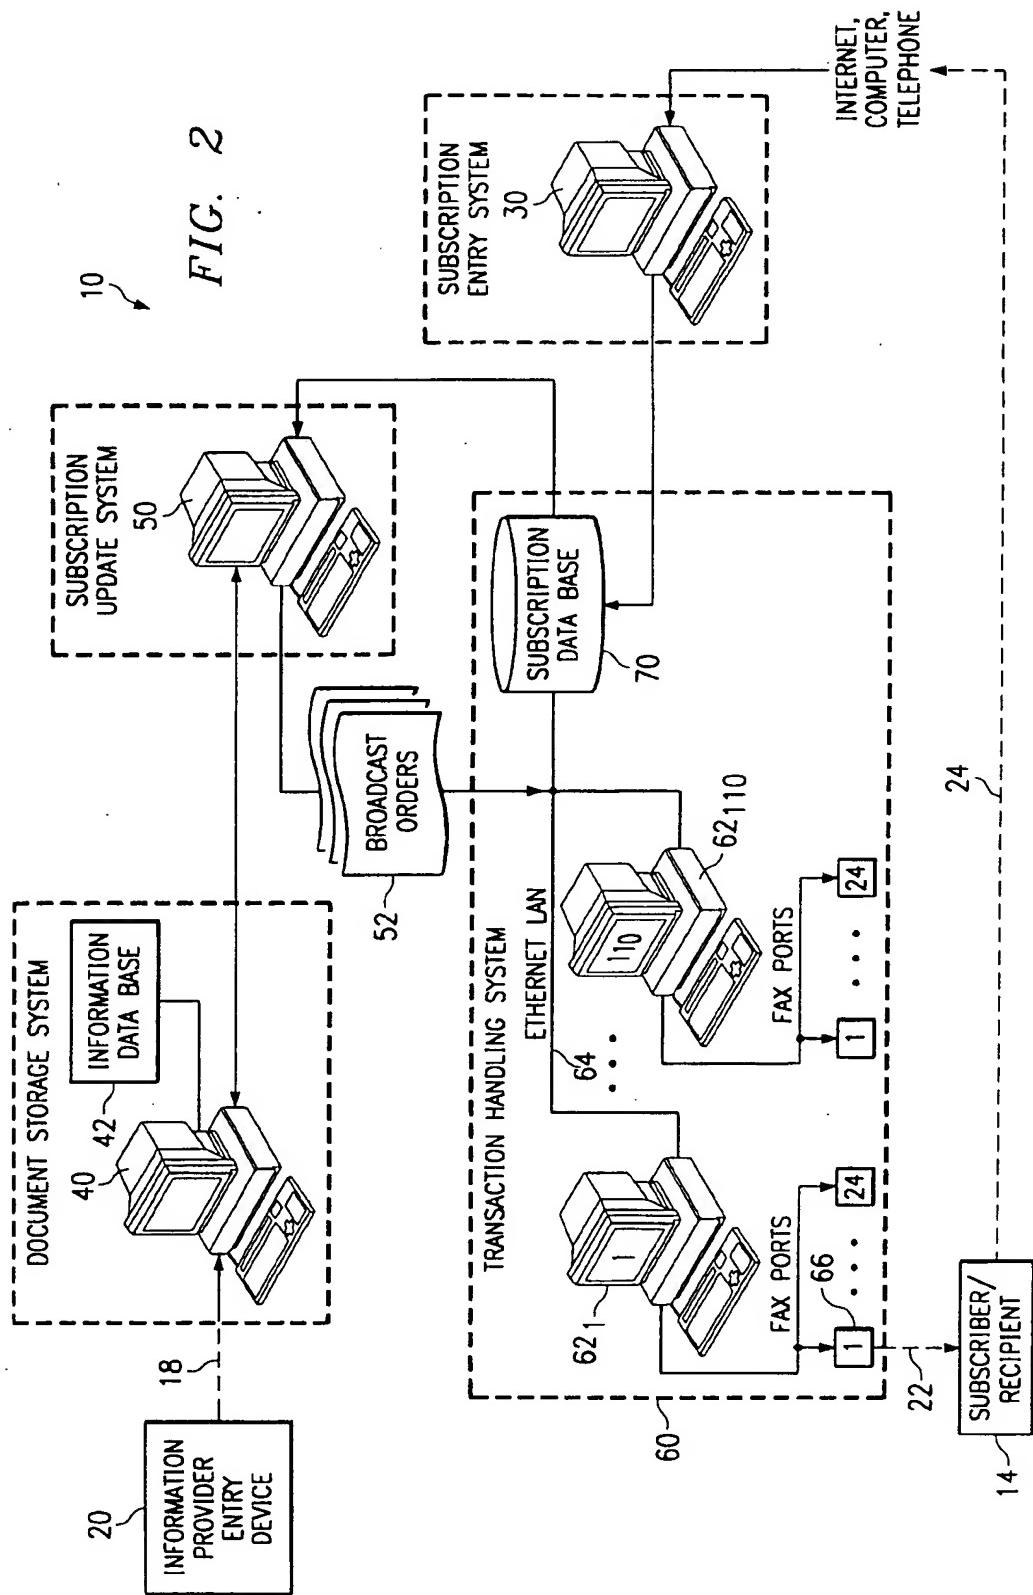


FIG. 3



Information Provider A Document Index

Title	Doc. Id.	No. of Pages	No. of Docs.	Fax	Time (min)
1.0 Document Index	1000	7	1		2
2.0 Tuberculosis Subject Matter	2000	51	3		21
2.1 Tuberculosis Diagnosis	2001	16	1		6
2.2 Tuberculosis Vaccination	2002	12	1		5
2.3 Tuberculosis Worldwide	2003	23	1		10
3.0 Polio Subject Matter	3000	383	12		160
3.1 Polio Control	3001	22	1		10
3.2 Polio Vaccination	3002	13	1		5
3.3 Polio in Asia	3003	38	1		15
	3004	12	1		5
		8	1		

80

FIG. 4

## SUBSCRIPTION ENTRY MODULE

1. Information for Each Subscriber
  - a) Information Provider ID
  - b) Fax Number
  - c) Password
  - d) Hours of Document Delivery
  - e) Document ID's or Groups Desired/Concealed for Document Delivery
  - f) Full Name - If within a subscriber group having one Fax
  - g) Update Information for a)-e) above
  - h) Historical Data
    - i) Document ID's Requested
    - ii) Total Documents Requested
    - iii) Number of Times the Index is Requested

FIG. 5

**INFORMATION DISTRIBUTION SYSTEM****TECHNICAL FIELD OF THE INVENTION**

The present invention relates in general to information distribution systems, and more particularly to systems for storing data, text and other types of information, and for distributing the information to subscribers in an efficient and timely manner without incurring either duplicated distributions or distribution to subscribers who are not interested in the particular information.

**BACKGROUND OF THE INVENTION**

Many aspects of a person's personal and business life depend on the availability of information from many sources. Indeed, the decisions that are made in a person's day-to-day life depend in a large part on the information that is not only available, but that which is actually received and considered by the individual. It can be appreciated that information that is available may nonetheless not be received or considered if the availability thereof is not sufficiently convenient. For example, libraries include thousands of sources of information, but the accessibility thereof is somewhat inconvenient, in that people must actually go to the libraries, look through the various indices of the authors, subject matter, etc., and then actually find the publication and study the same. The current use of the Internet and the associated mass of information has been widely accepted and used, primarily because of its ease of accessibility.

As noted above, the effectiveness of information depends on the speed, efficiency and reliability thereof. Timely accessibility and delivery increases the value of the information. An important technology that is widely used to facilitate the speed and delivery of the information, especially documents, is the facsimile (fax) mode of transmission. Fax delivery systems of the type that provide documentary information are available in two general categories. First, "fax broadcast" systems are utilized for delivering the same document at the information provider's request to a number of recipients at approximately the same period of time. On the other hand, "fax-on-demand" systems are available where an information provider places documents on a system and recipients can make telephone calls to request the desired documents. Both such systems are generally successful, but have serious deficiencies and limitations. The fax broadcast system is inefficient because it transmits all documents to all recipients identified on a list, regardless of whether the recipients are actually interested in the contents of the documents. Also, serious concerns remain with respect to the fax broadcast system in that the transmittal of unsolicited facsimile documents may violate various governmental rules and regulations.

A major shortcoming of the fax-on-demand system is its inefficiency, in that it requires a request of the information each time the recipient desires such information. Moreover, the recipient has no way of knowing when the information source has been changed, modified or even deleted.

The deficiencies of both the fax broadcast and fax-on-demand systems can be appreciated from the following example.

A hypothetical document distribution system is provided by the Center for Disease Control. The system satisfies a need to distribute a wide variety of warnings, notices and memoranda about diseases, viruses, diagnoses, treatments, medicines, etc., by fax to a large target group of doctors and technicians. The documents can be distributed by either the traditional fax broadcast technique, or the fax-on-demand

technique. In view that there can be tens of thousands of medical facilities included in the target group, there exists many documents that are not relevant to each of the facilities. As such, the fax broadcast technique is inefficient and very expensive. On the other hand, the fax-on-demand technique is impractical because the recipients have no way of knowing when new information becomes available, in view that the number of documents available is extremely large.

- 10 From the foregoing, it can be seen that a need exists for a new document distribution technique that is more efficient as to the distribution system, as well as the recipient. Another need exists for a new information distribution technique that allows recipients to receive only the information to which an interest exists, and automatically receive such information in a timely manner. Yet another need exists for an information distribution technique which allows recipients the capability to easily change or modify document requests such that new subject matters are easily added, and that the distribution of current subject matters are easily discontinued.

**SUMMARY OF THE INVENTION**

In accordance with the preferred embodiment of the invention, an information distribution system and method of operation thereof is disclosed for reducing or overcoming the shortcomings and disadvantages of the prior art techniques. In a preferred form of the invention, there is provided an information database for storing vast amounts of information and for providing high speed access thereto. The information database is adapted for allowing new and updated information to be easily added thereto at any time. In one application of the invention, the information database stores picture and character images of documents, as supplied to the system by way of facsimile and otherwise. A subscription database is provided for storing the identity of each subscriber or recipient. In addition to storing the identity of each subscriber, there is stored the subject matter areas of interest to the subscribers, as well as other necessary information.

A processor is programmed to cyclically scan the information database to determine if any new documents have been added thereto. As to any new documents that have been added to the information distribution system, the processor then scans the subscription database to identify each subscriber that has an interest in the subject matter of the new document. When the subject matter of the new documents has been correlated to the subscribers having an interest therein, the system causes an automatic transmission of the new documents to each subscriber having an interest therein. With this arrangement, the subscribers do not receive information in which they have no interest, nor are the subscribers required to formally request the new information as a prerequisite for its delivery to the recipient.

In accordance with another feature of the invention, the information distribution system is adapted for receiving communications from subscribers via facsimile, telephone calls, the Internet, etc., in which new subscribers can initially be registered with the system, and which current subscribers can add new areas of interest, delete others, and otherwise modify various parameters by which the information is delivered to the recipient.

In the preferred form of the invention, the information stored in the database is of the documentary-page type, and the distribution is by way of a facsimile engine to the various subscribers.

## BRIEF DESCRIPTION OF THE DRAWINGS

Further features and advantages will become apparent from the following and more particular description of the preferred and other embodiments of the invention, as illustrated in the accompanying drawings in which like reference characters generally refer to the same elements or parts throughout the views, and in which:

FIG. 1 is a generalized block diagram of the information distribution system, as electronically coupled to information providers and recipients who are registered as subscribers to the information delivery system;

FIG. 2 is a block diagram of the information distribution system constructed according to a preferred embodiment of the invention;

FIG. 3 is a listing of the data and parameters required in the preferred embodiment for each information provider as well as for the data for adding information to the information database;

FIG. 4 is illustrative of an example document index and the various subject matters provided by an information provider;

FIG. 5 is a listing of the types of parameters and data that are utilized for registering a subscriber with the system;

FIG. 6 is a listing of the functions carried out by a subscription update software module that cyclically scans the database for the appearance of new documents; and

FIG. 7 is a listing of the functions carried out in accordance with a document delivery software module for transmitting the information to a subscriber having an interest in the same.

## DETAILED DESCRIPTION OF THE INVENTION

With reference to FIG. 1 of the drawings, there is illustrated in block diagram form the connection of an information distribution system 10 between plural information providers 12<sub>a</sub>-12<sub>z</sub> and plural subscribers 14<sub>1</sub>-14<sub>n</sub>. The term "subscribers" is used synonymously with the term "recipients". When the communications between the information distribution system 10 and the subscribers 14 is by way of facsimile, such as is the case in the preferred embodiment of the invention, the public telephone system 16 provides the switched communications connection therebetween. As will be described more fully below, the subscribers 14 automatically receive new information of interest that has been input to the information distribution system 10. In addition, provisions are made for allowing any new subscriber to communicate with the information distribution system 10 to register as a subscriber for receiving information in different selected subject matter areas.

The information providers 12 can supply information to the distribution system 10 by way of facsimile and through many other mediums. Of course, the interface between the information providers 12 and the distribution system 10 would include suitable circuits for converting the input information to appropriate digital signals for storing on high speed mass storage devices, such as multiple disk drives. The information providers 12<sub>a</sub>-12<sub>z</sub> may each be involved in entirely different subject areas, all of which is stored in the information distribution system 10 and distributed in a timely manner to the various subscribers 16. For example, one information provider 12<sub>a</sub> may be involved in the control of diseases and may wish to have such information available only to those subscribers who are interested in the particular subject matter area. Another provider 12<sub>z</sub> may be involved

in automotive maintenance, which information will again be stored on the information distribution system 10 and made available to yet other subscribers, such as mechanics. A host of other diverse subject matter areas can be provided by other information providers, and stored in the system 10 so as to be available to the subscribers who have an interest in the same.

In a practical application, the information providers 12 may additionally provide goods and/or services to the subscribers outside of the system 10, while using the information as an incentive for doing business with the respective information providers 12. To that end, it is anticipated that the subscribers 14 will not have to pay a service fee for obtaining the information from the system 10. Rather, each time a new document or information is transmitted from the system 10 to a subscriber 14, the respective information provider 12 can be billed accordingly. Many other charge and payment schemes can be utilized by those skilled in the art. With the utilization of the invention for the distribution of information, the sales of goods and services by the information providers 12 can be facilitated. Indeed, new product information, product brochures, press releases concerning products, etc., can be made readily available to the subscribers, via the information distribution system 10.

With reference now to FIG. 2, there is shown a more detailed block diagram of the information distribution system 10, as coupled to one information provider entry device 20 by way of the communication line 18. It should be understood that many other information providers 12 would typically be coupled to the system 10. In the preferred embodiment of the invention, the information provider entry device 20 comprises a facsimile machine for transmitting standard facsimile format signals, such as the T30 fax protocol, on line 18 to the information distribution system 10. A subscriber 14, or recipient, receives a facsimile transmission by way of the communication line shown by the broken line 22. Moreover, the subscriber 14 can communicate with the information distribution system 10 by a communication medium 24 which may include the Internet, a computer line, telephone line, etc. In the preferred form of the invention, new subscribers can communicate with the information distribution system 10 on various communication mediums 24 when initially registering with the system to receive facsimile transmissions concerning specific subject matter areas.

The major functional subsystems of the information distribution system 10 include a subscription entry system 30 in which registration information from new subscribers is entered into a subscription database 70; a document storage system 40 for storing text and images input by the information providers 12 into an information database 42; a subscription update system 50 for periodically scanning the information database 42 for the occurrence of new information; and a transaction handling system 60 which comprises the facsimile engine for transmitting new documents to the subscribers interested in the subject matter thereof. The transaction handling system 60 is of the same type and construction as used in the conventional facsimile broadcast system. In the preferred embodiment of the invention, up to 110 personal computers, one shown as reference number 62, are connected together by an Ethernet local area network 64. Each personal computer 62 is, in turn, associated with up to 24 facsimile ports, one shown as reference numeral 66. With 110 personal computers, each driving 24 facsimile ports, a total of 2640 facsimile ports can be simultaneously driven to transmit information to the respective subscribers 14. The local area network 64 also couples the personal computers

62 of the facsimile engine to the subscription database 70. The subscription database 70 stores all the data that is uniquely related to each subscriber 14. As will be set forth more fully below, the personal computer 30 in this subscription entry system receives all of the data from the subscribers, in response to prompts and the like, and stores such data in the subscription database 70.

As noted in FIG. 2, the subscription update system 50 has access to the subscription database 70. Moreover, the subscription update system 50 also has access to the information database 42. Periodically, the subscription update system 50 reads the information database 42 to determine all new or revised documents that have been added thereto. Based on the subject matter area of each new or revised document stored in the information database 42, the subscription update system 50 accesses the subscription database 70 to determine if any of the subscribers has registered an interest in one or more of the documents. On a match between the subject matter of a new or revised document and a corresponding subscriber interested in the same, the subscription update system 50 transmits broadcast orders 52 to the transaction handling system 60 for causing a facsimile transmission of each new or revised document to the respective subscribers having an interest in the subject matter thereof. Any document that has been revised in any manner is considered as a "new" document.

The operation of the system is as set forth below. An information provider 12 utilizes an entry device 20, such as a facsimile machine, for accessing the information distribution system 10 and particularly the document storage system 40. When employing facsimile transmissions, a telephone line 18 can be utilized by dialing one or more facsimile numbers associated with the document storage system 40. Alternatively, information providers 12 can access the information distribution system 10 by way of the Internet to also convey document data, wave files, voice information, etc. Numerous other communication mediums can be utilized for transmitting the different types of information to the information distribution system 10.

The information provider 12 can classify and prioritize the subject matter of the documents, as well as the information data in numerous different ways. Various parameters can also be transmitted according to a predefined format, in association with the information data. Such parameters may include the title of the document, the document identification, the number of pages, information provider identification number, a password, the expiration date and/or effective date of the document, the approximate facsimile time of the document, priority, and other possible helpful information. Such data is stored in the information database 42 in the manner transmitted by the information provider entry device 20. Importantly, the information and all associated parameters are stored in the information database 42, in the manner substantially identical as transmitted by the information provider 12.

The sophistication or format in which the data is stored in the database 42 is primarily a function of versatility or efficiency by which the data can be searched, categorized and found. In the preferred embodiment of the invention, each document is stored in association with the following parameters, shown more fully in FIG. 3. The record stored in the data base for each document includes a number of fields, each storing the various parameters. An identification code unique to each information provider 12 is transmitted by the entry device 20. With such a code, the information distribution system 10 can maintain separate records of each information provider 12. The billing address for the infor-

mation provider is also stored so that at the end of each billing period, the information provider is billed, based on the number of times subscribers requested each document. A password associated with the information provider may also be required. Lastly, the data base may store a subscriber list which identifies each of the subscribers authorized by the information provider to subscribe to and obtain facsimile transmissions of the provider's documents. Next, the transmission of each document by an information provider 12 to the system 10 is preferably accompanied by a password that verifies a valid transmission by the information provider to be stored in the associated file. This password is verified by the information distribution system 10. This prevents hackers and other unauthorized persons from intentionally or inadvertently transmitting files for storage in the wrong data base files.

Each document registered with the information distribution system 10 includes a document identification number or code that uniquely identifies the document. This ID number is assigned to each document by the system 10. The title of the document, the number of pages, the approximate fax time for delivery, the activation date, the expiration date and a priority (if any) are also transmitted with the document to the system 10. If the information provider 12 has more than one group of subject matter areas, then each group is identified by a unique identification number. This allows the documents of each group to be maintained in a separate file so that the subject matter areas can also be maintained and accessed separately. This group identification code also allows a subscriber to simply request receipt of a group of documents in the same subject matter area without requesting the documents individually.

In addition to the document information stored in the database 42, the information provider 12 may also provide a document index 80, an example of which is shown in FIG. 4. Here, the document index 80 is preferably a document itself that is prepared by the information provider 12 and transmitted upon request to a subscriber 14. The document index 80 lists the various groups of all of the documents provided by the information provider 12. In the document index 80 shown in FIG. 4, it is noted that the index itself is identified as document identification 1000, having seven pages, comprising a single document, and requiring about two minutes of facsimile time for transmission thereof. The index 80 also shows that a first subject matter group 2.0 pertains to tuberculosis, while a second subject matter group 3.0 relates to a polio subject matter. The specific documents within the tuberculosis subject matter group 2.0 include tuberculosis diagnosis document identification 2001, tuberculosis vaccination document identification 2002 and tuberculosis worldwide document identification 2003. Each of the three individual documents 2.1-2.3 are associated in the index 80 with the number of pages, the number of documents and the approximate facsimile time. If all documents in the tuberculosis subject matter area are desired, then the group 2.0 can simply be requested, in which event all 51 pages of the three documents will be transmitted, which facsimile transmission time will be about 21 minutes.

The second subject matter group shown in the index 80 of FIG. 4 illustrates a polio group 3.0, having a total of 383 pages comprising 12 documents, with an approximate facsimile time of 160 minutes. Listed under the polio subject matter group 3.0 are the numerous individual documents, only a portion of which are shown. While the index 80 of FIG. 4 is illustrative of the many variations that are possible, those skilled in the art may desire to provide a much more expansive index, a simplified index, or no index at all. As

will be set forth below in more detail, when a subscriber or recipient 14 initially registers with the information distribution system 10, the recipient may request that the document index 80 be transmitted with the transmission of each new document. Alternatively, the recipient 14 may request that an index 80 of the particular information provider 12 only be transmitted once a week, once a month, etc., irrespective of the number of documents transmitted therebetween. Each information provider 20<sub>a</sub>-20<sub>x</sub> can add new documents or information to the information database 42 at any time whatsoever.

The subscription entry system 30 is provided so that current subscribers can access the system and request a facsimile transmission of any document, including the document index 80, provided by the respective information providers 12. As noted above, it is contemplated that each information provider will supply to the information distribution system 10 a listing of all of the subscribers who can request documents provided by the respective information provider 20. This prevents unauthorized subscribers from requesting the transmittal of unnecessary documents for the sole purpose of running up the bill charged by the information distribution system 10 to the providers 12. In order for a subscriber/recipient 14 to request a document, the document identification code can be transmitted on a medium 24 according to a specified format, together with the subscriber identification code and other relevant information, as may be required. Alternatively, the subscriber 14 can request documents by way of the Internet and the use of a checklist, by telephone or other suitable means.

Returning to FIG. 3, there is written in a field of the data base record of each document a calendar date and time of day in which the document was registered with the system 10. By date and time stamping each document, it can be readily ascertained whether other documents having the same title are newer versions.

In addition to the request of documents, the subscription entry system 30 is also programmed to allow new subscribers to register therewith, so as to automatically receive any or all new documents that have been stored in the information database 42 by the respective information provider 20. The registration on the information distribution system 10 by a new subscriber 14 may also be carried out by way of facsimile transmission, Internet, computer data, telephone, etc. When using a telephone to request documents or to be registered as a new subscriber, either operators or interactive communications via voice prompts may be employed. The relevant information to be transmitted from the new subscriber to the subscription entry system 30 is shown in FIG. 5, and comprises the facsimile number of the subscriber, a password assigned to the subscriber by the information provider 20, the group identification code pertaining to the subject matter of interest to the subscriber, and the hours between which new documents can be transmitted to the subscriber 14. Of course, in the event that the new subscriber 14 is associated with a group of recipients utilizing the same facsimile machine, then the name of the recipient may additionally be required. When registering by way of the Internet, the subscriber can access the web site of the information distribution system 10, whereby a checklist is provided to enter the information noted above. Further, the document index 80 provided by the information provider 12 may also be accessible, whereby the new subscriber 14 simply checks the relevant box beside each group, or each document, to indicate an interest therein. In accordance with an important feature of the invention, current subscribers can additionally access the subscription entry system 30 to

modify the associated registration information, such as when a new facsimile number is to be used, new areas of interest in either groups or documents are added, or deleted. If, for example, the recipient 14 intends to be absent for a period of time and desires that no documents are to be transmitted, then the check marks of all of the boxes relating to the groups or documents are deleted, whereby the recipient remains registered with the system, but no documents are delivered until the recipient again provides the relevant information as to the subject matter areas.

The subscriber registration data entered into the subscription entry system 30 is stored in the subscription database 70, as shown in FIG. 2. As noted in such figure, subscriber data stored in the subscription database 70 is available to the subscription update system 50.

According to an important feature of the invention, the major functions carried out by the subscription update system 50 are shown in FIG. 6. The update system 50 is programmed to operate on a cyclical basis, and preferably about once each hour, to access the information database 42 and identify each new document that has been entered into the information distribution system 10 during the immediate preceding hour. For each new document entered into the system 10 during such time period, a transaction record is generated. The transaction record includes the document ID and data related to the subject matter of the document. The subscription update system 50 searches through the entire information database 42 for each new document. As noted above, a transaction record is generated for each such document found. Then, the subscription update system 50 accesses the subscription database 70 to find a correlation between the subject matter of each new or revised document, and the subscribers registering an interest in such subject matters. If there is no correlation found between the subject matter of the new documents and at least one subscriber, then the corresponding transaction record is not transmitted. On the other hand, as to all of the subject matter areas of the new documents in which there is a correlated subscriber, the record of each new document is completed with the name and facsimile number of the subscriber. One transaction record for each new document in which there exists a subscriber having an interest in the subject matter thereof is transmitted in a serial manner as a broadcast order 52 shown in FIG. 2. The broadcast orders 52 are transmitted to the local area network 64 for distribution to the respective personal computers 62<sub>1</sub>-62<sub>110</sub> of the transaction handling system 60.

The image data of each new document is not part of the respective transaction record. Rather, when the transaction record is assigned to a particular personal computer 62 in the transaction handling system 60, such personal computer decodes the various fields of the transaction record. The outbound facsimile port then dials the facsimile number using the document ID as an index. Prior to the dialing of the facsimile number of the subscriber, the personal computer 62 retrieves the image data of the respective document from the information database 42. With such information, when the facsimile machine of the subscriber responds by going off-hook, the document image is delivered.

The format of delivery can be of different variations, including a face or header sheet having the recipient's name, and any other information desirable. In the event the subscriber's facsimile machine is unavailable for receiving the transmission, such as being in use, out of paper, etc., the personal computer 62 will attempt a retransmission after a predefined delay period. After a predetermined number of attempts to deliver the document image, the personal com-

puter 62 of the transaction handling system 60 will flag the record as undelivered. Records flagged as such can be queued so as to reattempt transmission during the next cyclical period.

The personal computer 62 of the transaction handling system 60 is also programmed to maintain historical records and other administrative records for billing purposes. To that end, a billing record is generated to indicate for each information provider 12 the cumulative number of each document delivered, the recipients to whom the document was delivered and the time required for facsimile delivery. This information is employed and gathered on a monthly basis so that each information provider 12 can be billed accordingly. Moreover, such historical data is made available to each information provider 12 for "housekeeping purposes", so that documents never or seldom requested can be deleted from the system.

From the foregoing, it can be seen that the information delivery system facilitates the dispersal of information in an efficient manner. Unlike the traditional facsimile broadcast technique, the documents are delivered only to those recipients who have an interest in the subject matter thereof. Unlike the fax-on-demand technique, new documents entered into the system are automatically delivered to the recipients interested in the same, without such recipients having to make specific requests for each such delivery. Shown and described above are the fundamental novel features of the invention, as applied to the preferred embodiment. It will be understood that various omissions, substitutions and changes in form and detail of the invention as described herein may be made by those skilled in the art, without departing from the true spirit and scope of the invention as defined by the appended claims. Therefore, it is the intention that the invention be limited only by the scope of the following claims.

While the preferred embodiment of the invention has been described in connection with plural programmed processors or personal computers, the entire system can be controlled centrally by the use of a programmed mainframe computer, or the like. As such, the various software modules would be incorporated together, or executed individually by the mainframe computer.

What is claimed is:

1. An information distribution system, comprising:  
an information database for storing information, and to  
which new information can be added;  
a subscription database for storing an identity of each of  
a plurality of subscribers, and for storing an area of  
subject matter interest of each said subscriber;  
a processor programmed to periodically cause accessing  
of said information database to identify the new informa-  
tion that has been added thereto subsequent to a  
prior periodic accessing of said information database;  
and  
said processor being programmed to cause accessing of  
said subscription database to identify each subscriber  
having an interest in the new identified information,  
and said processor causing an automatic transmittal of  
the new identified information to respective subscribers  
having an interest therein, whereby the respective sub-  
scribers receive said new identified information with-  
out having specifically requested the same.
2. The information distribution system of claim 1,  
wherein said information database is adapted for storing  
images of documentary information, and each document is  
associated with an identification code.

3. The information distribution system of claim 1,  
wherein said new information is associated with a code  
relating to a subject matter of the information, whereby  
different subject matter codes are employed with different  
information that relate to different subject matter areas.

4. The information distribution system of claim 1,  
wherein said subscription database includes in association  
with each subscriber identity, a facsimile number, and  
wherein said transmittal of said information is by way of  
facsimile.

5. The information distribution system of claim 1, further  
including an index stored in a database, said index identi-  
fying different groups of said information.

6. The information distribution system of claim 5,  
wherein a plurality of documents having a similar subject  
matter are stored as a group such that retrieval of the group  
retrieves all the documents belonging to the group.

7. The information distribution system of claim 6,  
wherein in response to a request by a subscriber of a group,  
each document belonging to the group is transmitted to the  
subscriber.

8. The information distribution system of claim 1, further  
including subscription entry apparatus, said apparatus  
including means for allowing subscribers to communicate  
with said information distribution system to input one or  
more subject matter interests, and means for associating the  
subject matter interest with the subscriber.

9. The information distribution system of claim 1,  
wherein new information entered for storage in the system  
is assigned a subject matter thereto and assigned an infor-  
mation identification code, and at periodic intervals said  
processor identifies the new information entered into the  
system since a prior time interval, correlates the subject  
matter thereof to subscribers who have an interest in such  
subject matter, and causes transmission of the new informa-  
tion to only those subscribers who have an interest therein.

10. A method of distributing information, comprising the  
steps of:

providing information to an information database from a  
plurality of information providers, each information  
provider being identified by a unique code;

providing by each said information provider an identity of  
each subscriber who can request the information pro-  
vided to said information database by the respective  
information provider;

storing the information supplied by the information pro-  
vider in the information database in association with a  
subject matter area thereof;

storing the identity of each of a plurality of subscribers in  
a subscriber database, and storing a subject matter  
interest of each said subscriber;

accessing the information database to determine subject  
matter areas and accessing the subscriber database to  
identify the subscribers having an interest in the subject  
matter area; and

transmitting the information to the respective subscribers  
having an interest in the subject matter area thereof.

11. The method of claim 10, further including periodically  
accessing the information database to determine new infor-  
mation added to the information database, and then identi-  
fying the subject matter area of the new information.

12. The method of claim 11, further including transmitting  
information to a subscriber that has been specifically  
requested, where the information is not new information.

13. The method of claim 10, further including storing an  
index in said information database, said index defining the

**11**

subject matter areas and an identity of the information corresponding to the subject matter areas.

14. The method of claim 13, further including allowing subscribers to subscribe to the index.

15. The method of claim 13, further including storing a different index for each of said plurality of information providers providing the information for storage in the information database.

16. The method of claim 10, wherein said information comprises documents, and further including storing in a common group different documents having a similar subject matter area.

17. The method of claim 16, further including allowing a subscriber to request a group subject matter area, whereupon each document assigned to the group is transmitted to the subscriber.

18. The method of claim 10, further including transmitting the information to the subscriber by way of facsimile.

19. The method of claim 10, further including allowing by said subscriber access to the information stored in said information database without charge, and charging an access cost of subscribers to said information providers for access by each subscriber to the respective information.

20. A document distribution system, comprising:

data entry apparatus including a processor and a document database for receiving and storing data corre-

**12**

sponding to document images, said processor being programmed to store in association for each document image a document identification code, a document title, a number of pages and a subject matter of the document;

subscription entry apparatus including a processor and a subscription database for receiving data from a plurality of subscribers, said processor being programmed to store in association with each subscriber a destination transmission number, a unique password, a subject matter area of interest, and a window of time in which information can be transmitted to the user; and

subscription update apparatus including a processor programmed to periodically scan said document database and scan for the existence of any new document entered as data into said system, and programmed to correlate a subject matter of each new document to a corresponding subject matter area of content to one or more of said subscribers, and programmed to cause a transmission of each new document to each subscriber in which the document subject matter matches the subject matter area of interest of the subscriber.

\* \* \* \* \*



US006044372A

**United States Patent [19]****Rothfus et al.****Patent Number: 6,044,372****Date of Patent: Mar. 28, 2000**

[54] **METHOD AND APPARATUS FOR PUBLISHING INFORMATION TO A COMMUNICATIONS NETWORK AND ENABLING SUBSCRIPTIONS TO SUCH INFORMATION**

[75] Inventors: **Eric J. Rothfus; Patrick J. Hill**, both of Austin, Tex.

[73] Assignee: **Dazel Corporation**, Austin, Tex.

[21] Appl. No.: **08/896,783**

[22] Filed: **Jul. 18, 1997**

[51] Int. Cl.<sup>7</sup> ..... **C06F 15/173**

[52] U.S. Cl. ..... **707/10; 707/104; 395/200.31; 395/200.42; 395/200.43; 395/200.53**

[58] Field of Search ..... **707/10, 104; 395/200.31, 395/200.42, 200.43, 200.53**

[56] **References Cited**

**U.S. PATENT DOCUMENTS**

5,732,219	3/1998	Blumer et al.	395/200.57
5,737,599	4/1998	Rowe et al.	707/104
5,737,619	4/1998	Judson	707/500
5,742,762	4/1998	Scholl et al.	395/200.3
5,745,754	4/1998	Lagarde et al.	395/200.33
5,752,022	5/1998	Chiu et al.	707/10
5,761,673	6/1998	Bookman et al.	707/104
5,764,906	6/1998	Edelstein et al.	395/200.49

5,778,367 7/1998 Wesinger et al. .... 707/10  
 5,793,966 8/1998 Amstein et al. .... 395/200.33  
 5,796,952 8/1998 Davis et al. .... 395/200.54

*Primary Examiner—Paul R. Lintz*

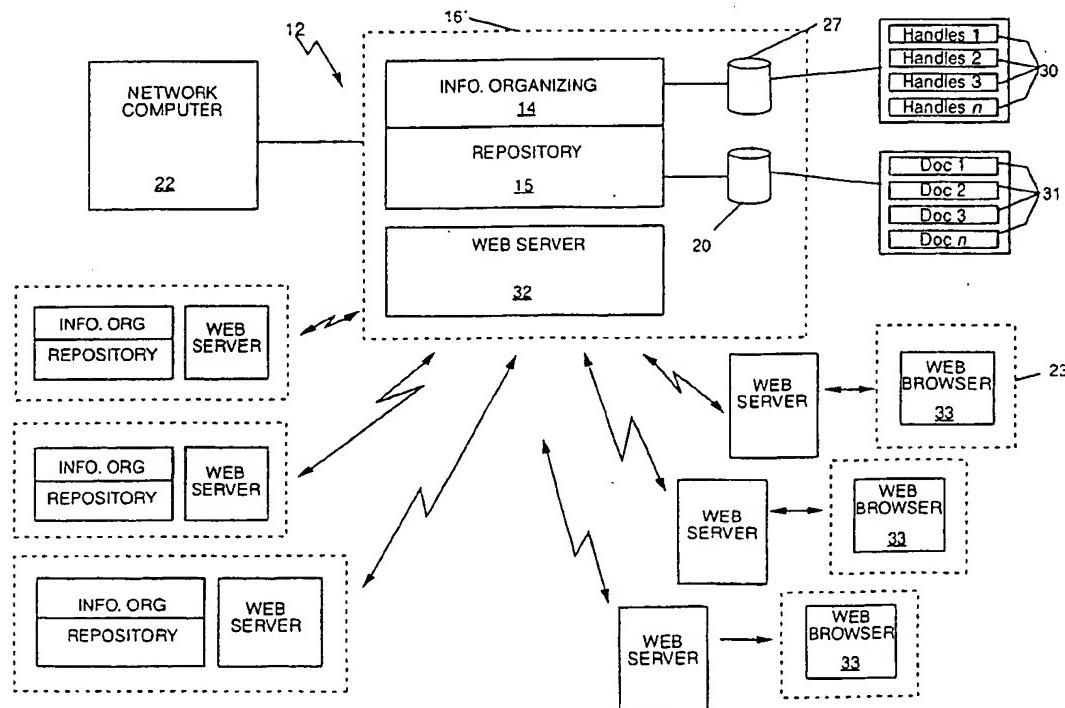
*Assistant Examiner—Jean R Homere*

*Attorney, Agent, or Firm—Russell D. Culbertson; Shaffer & Culbertson, LLP*

[57] **ABSTRACT**

An information element to be published is first associated with a first network addressable location of a communications network. A repository arrangement programmed in a computer receives the information element to be published and stores the information element at a second network addressable location. The computer implementing the repository arrangement is connected to the communications network. When the information element is stored by the repository arrangement, a handle production arrangement programmed in the computer produces a repository handle for the information element. The repository handle includes identifying information relating to the information element to be published and information from which a network address for the second network addressable location may be produced. An information organizing arrangement programmed in the computer retrieves the repository handle for the information element to be published and stores the repository handle under a first network addressable location.

33 Claims, 9 Drawing Sheets



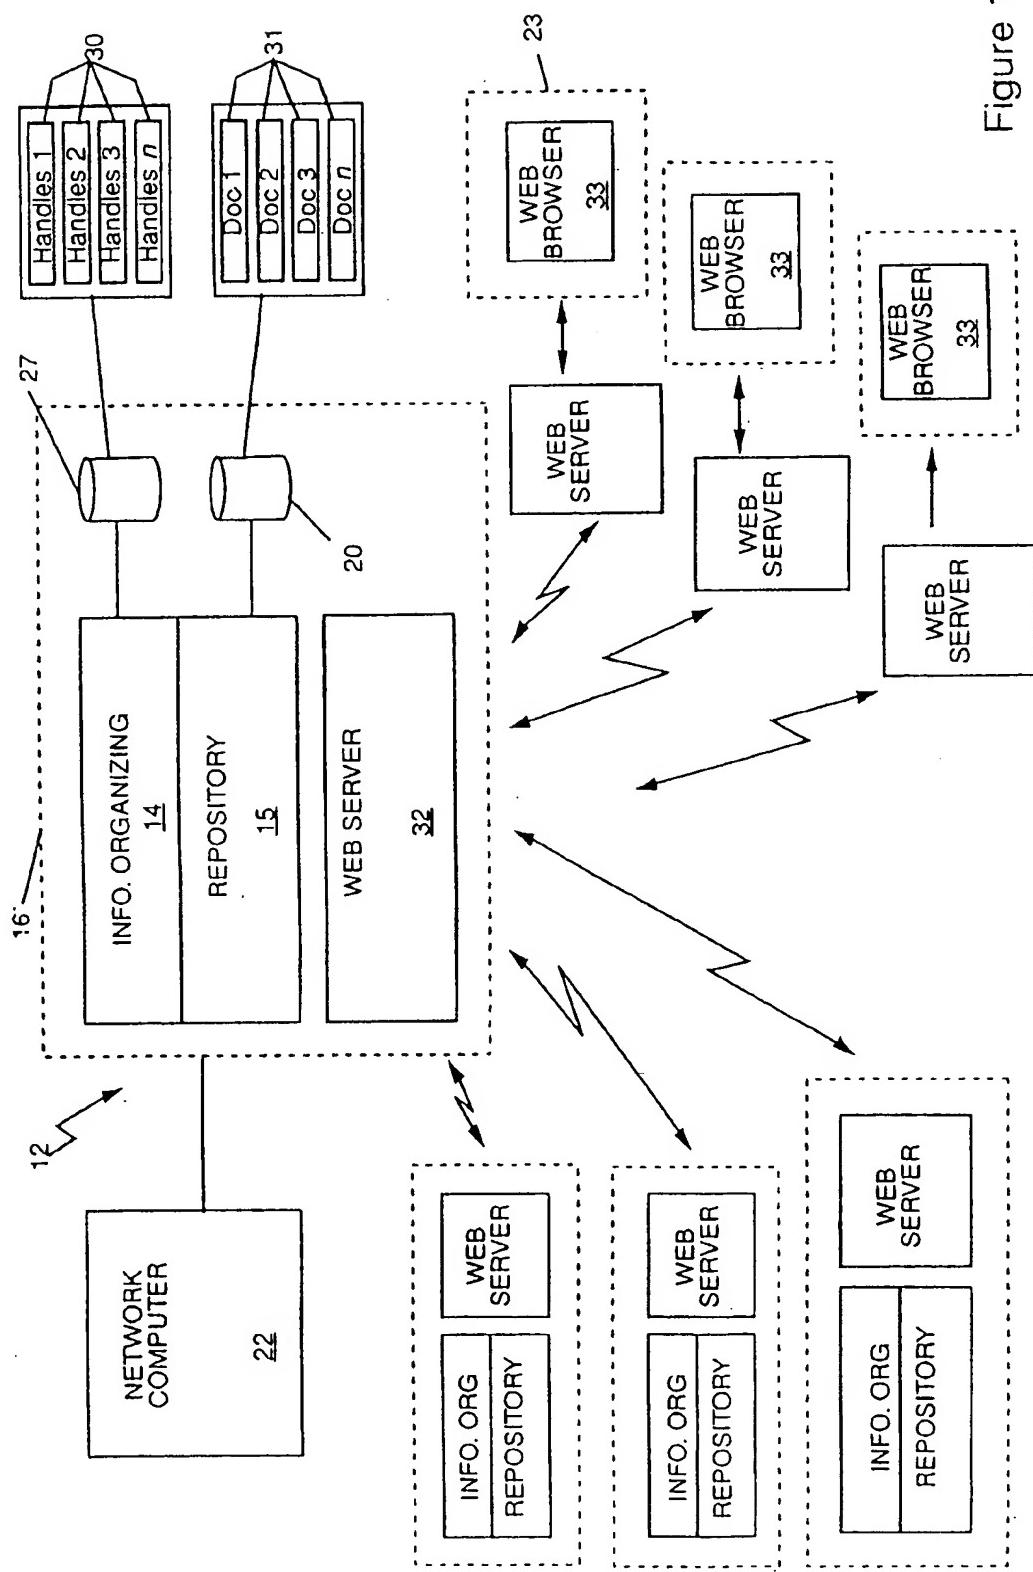


Figure 1

Figure 1A

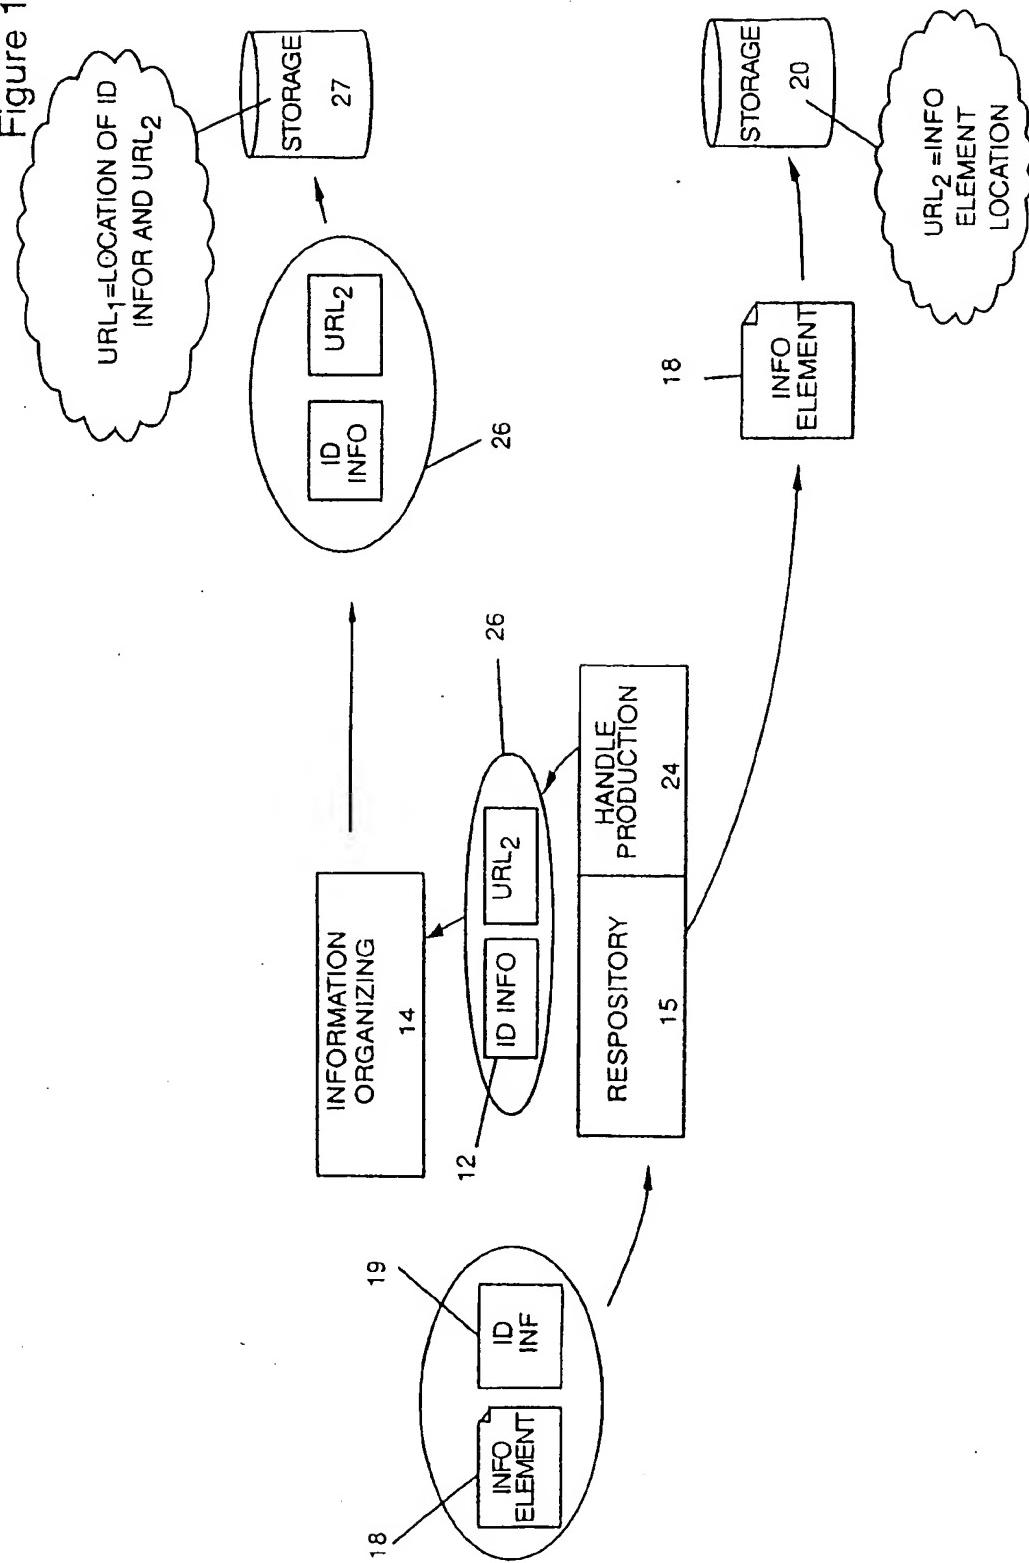


Figure 2A

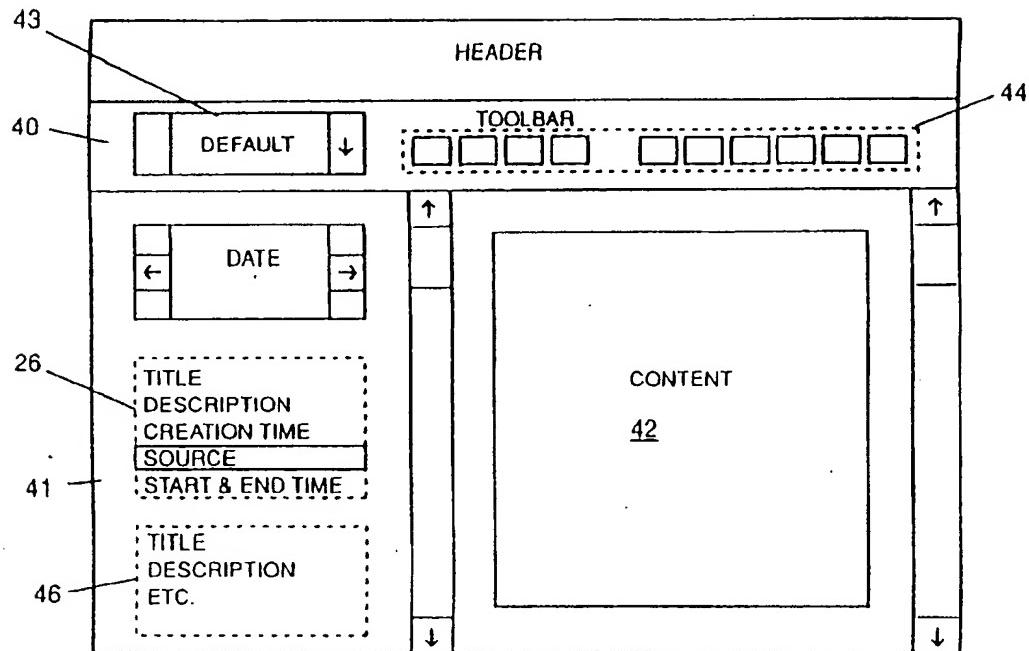


Figure 2B

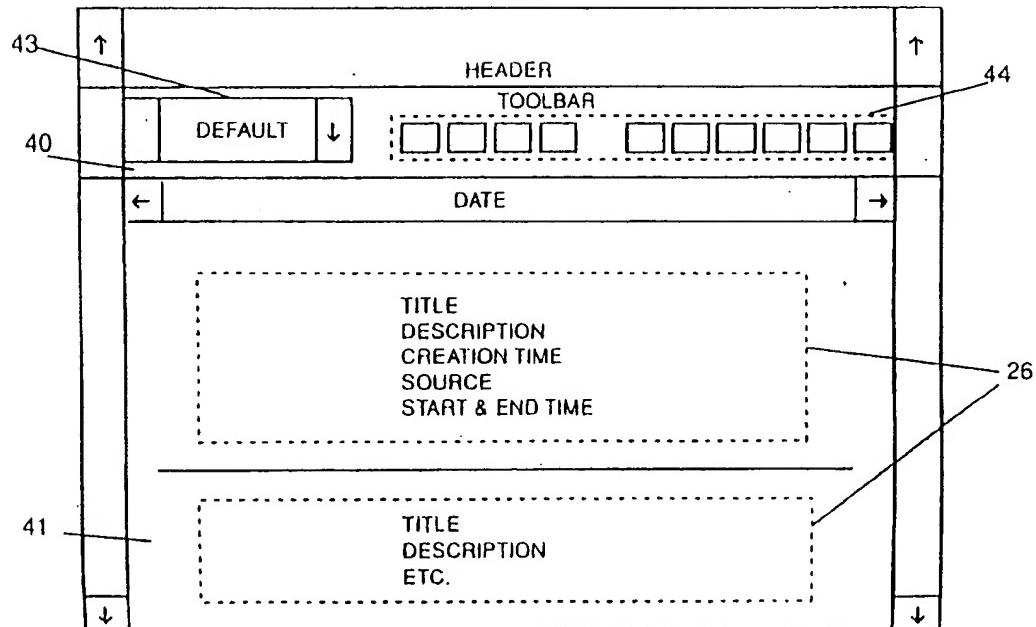


Figure 3

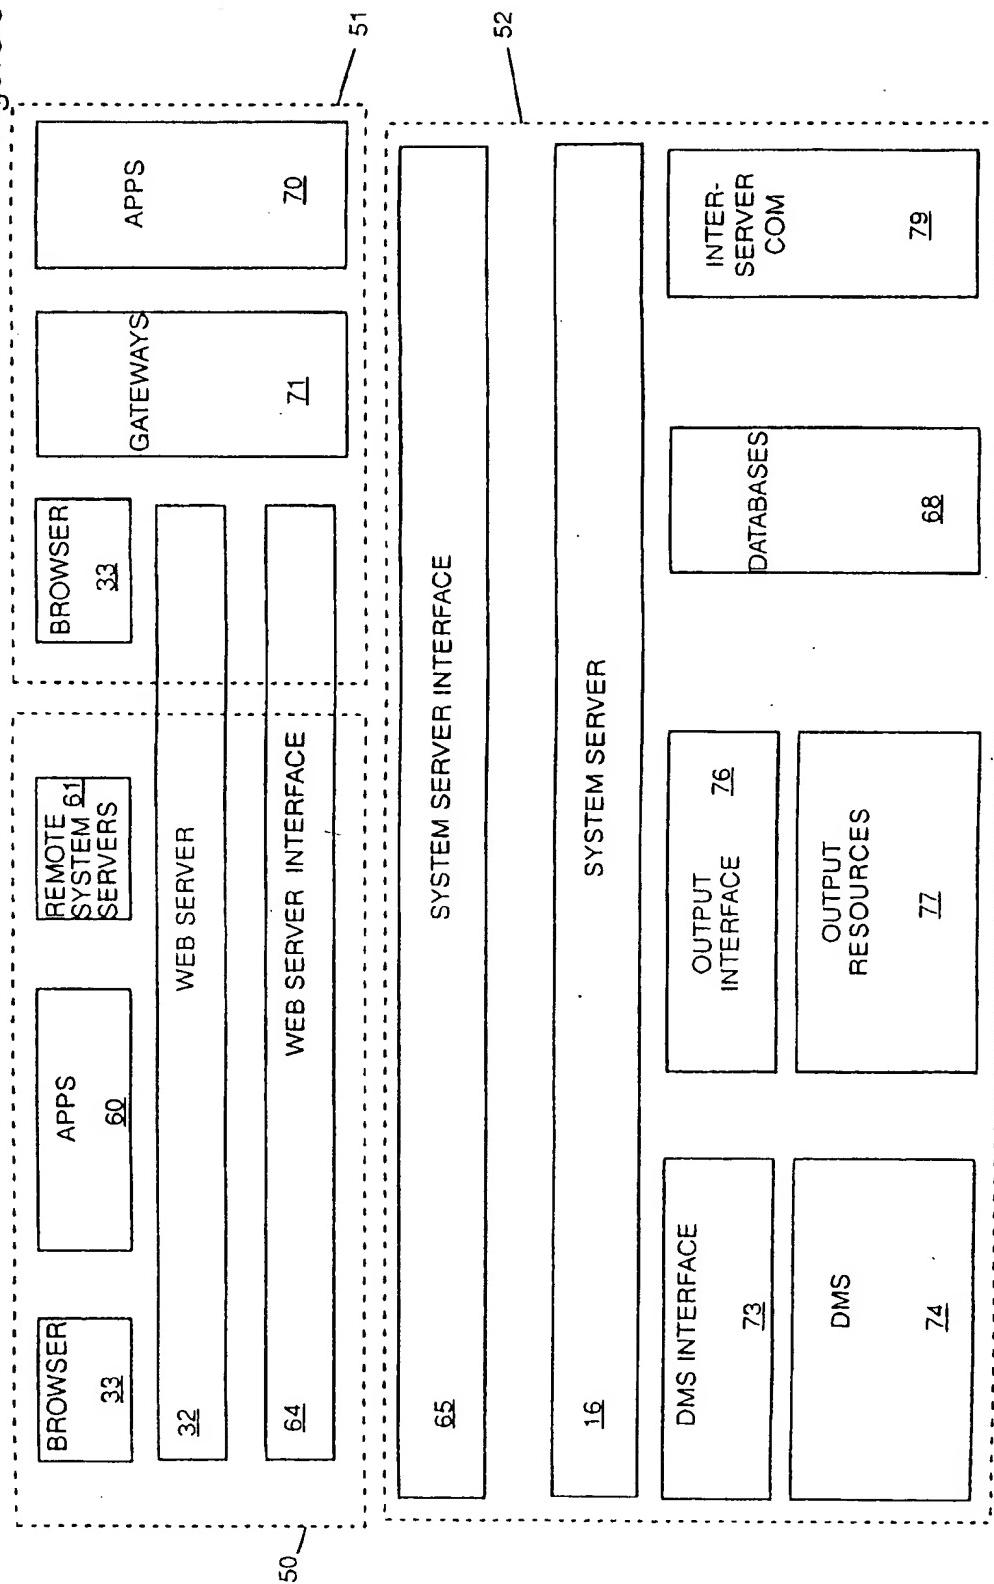


Figure 4

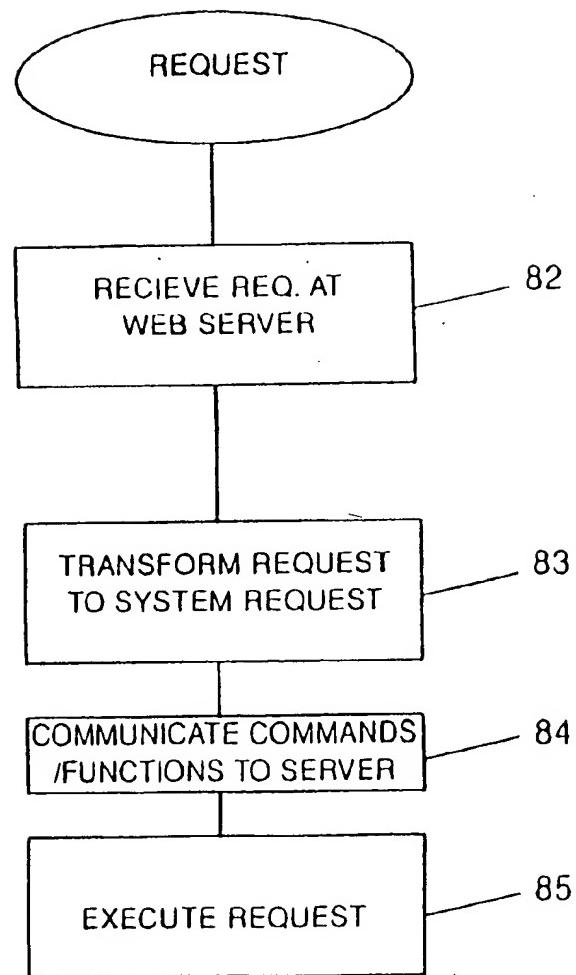


Figure 5

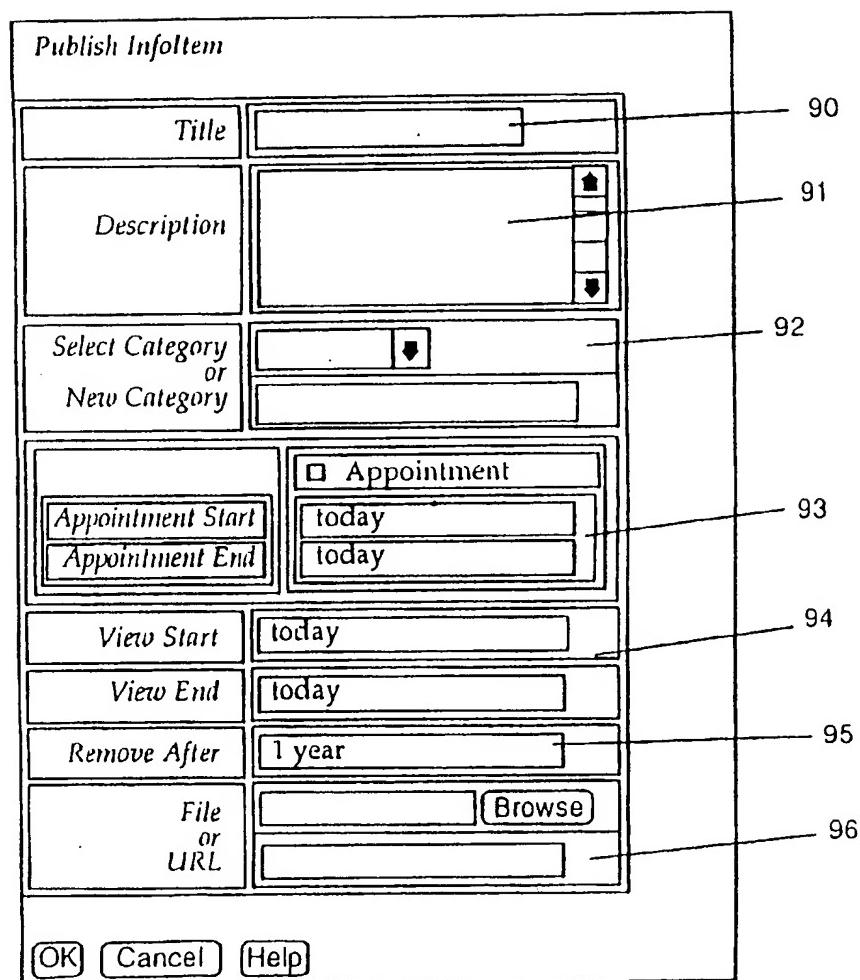


Figure 6

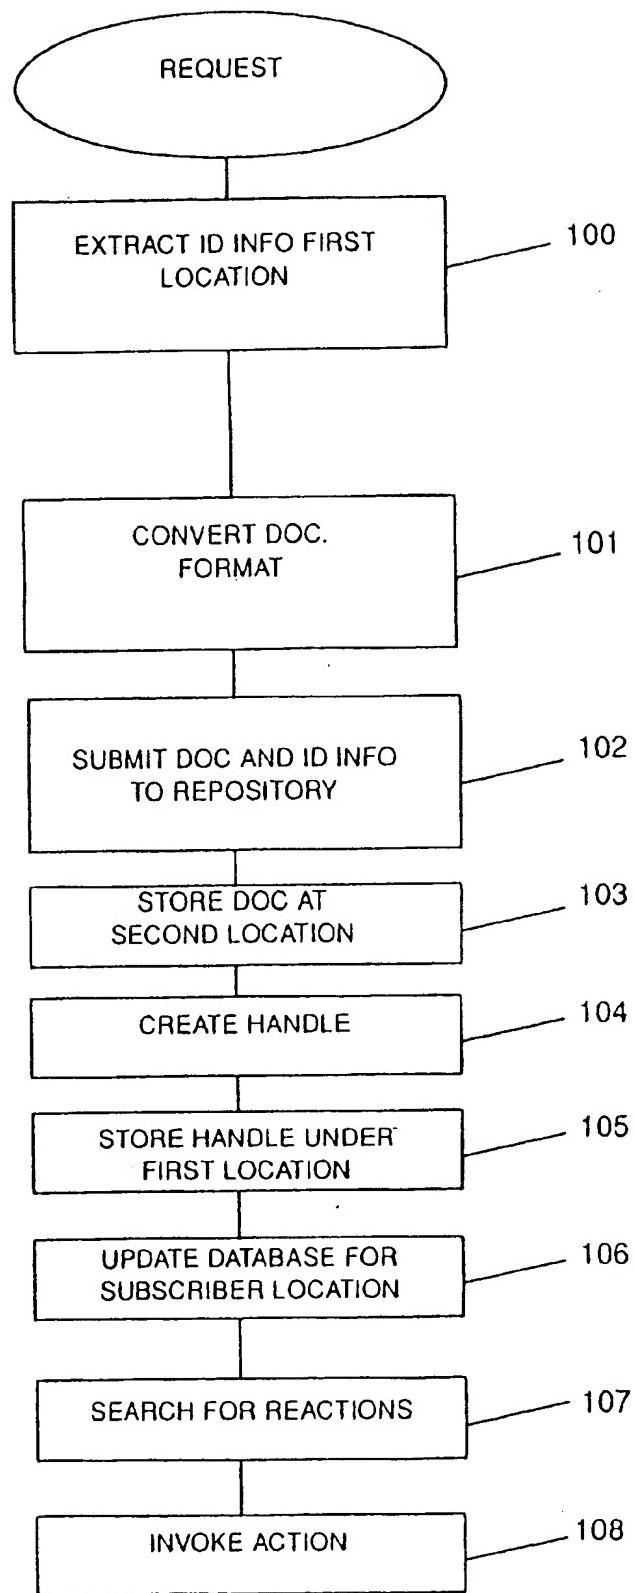


Figure 7A

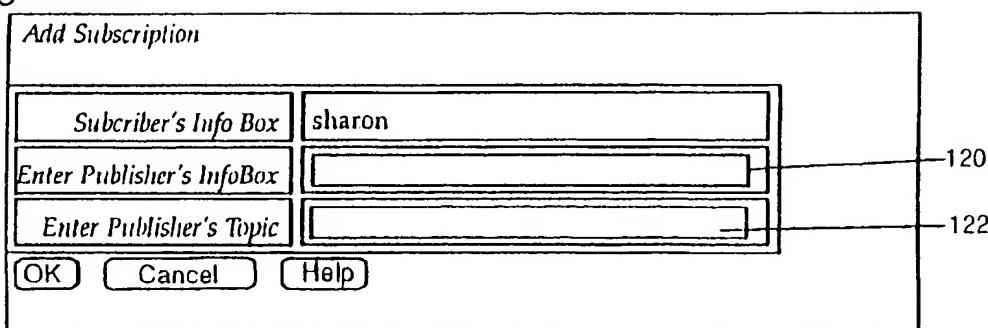


Figure 7B

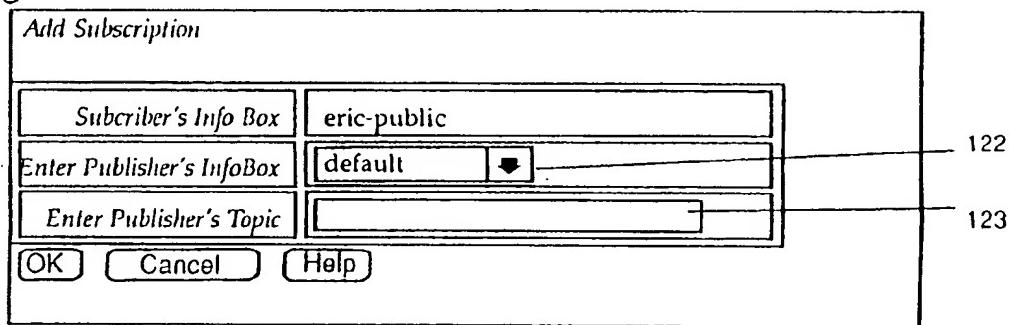


Figure 8

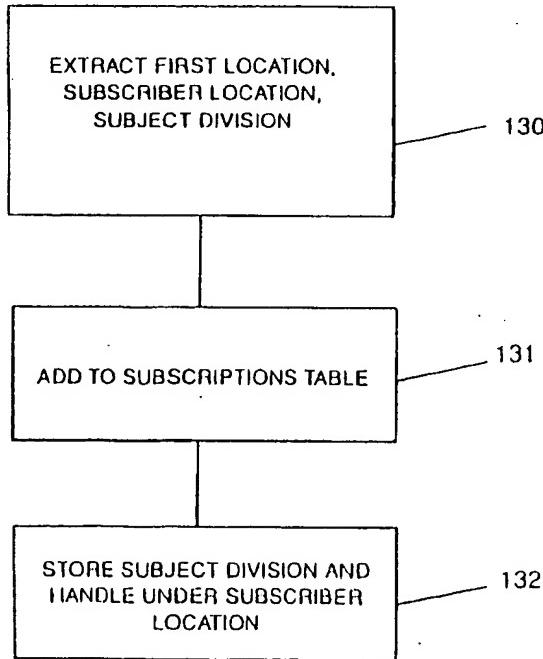


Figure 9A

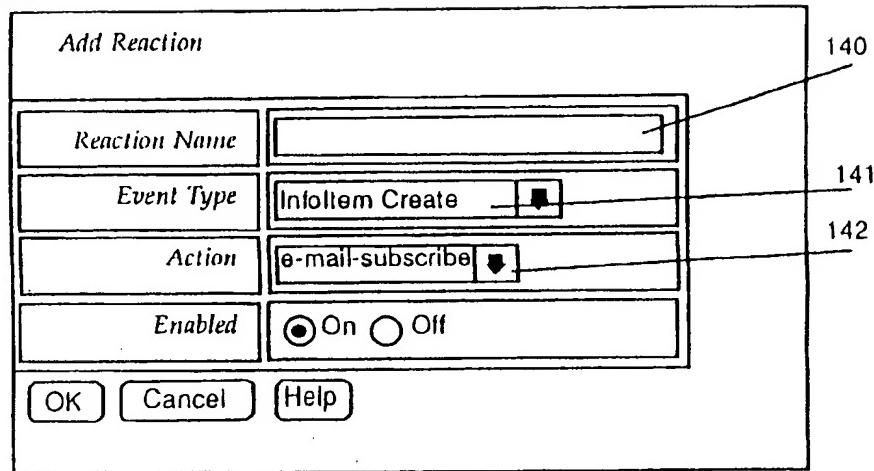
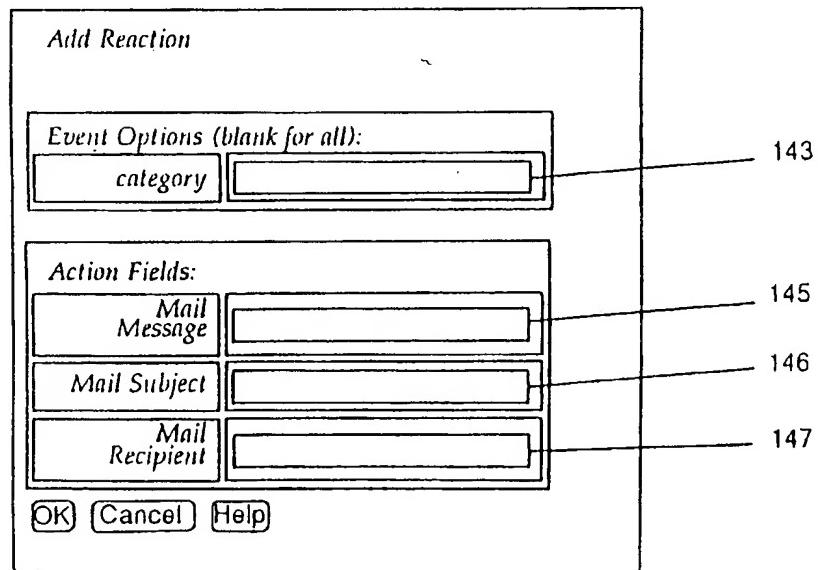


Figure 9B



**METHOD AND APPARATUS FOR  
PUBLISHING INFORMATION TO A  
COMMUNICATIONS NETWORK AND  
ENABLING SUBSCRIPTIONS TO SUCH  
INFORMATION**

**BACKGROUND OF THE INVENTION**

This invention relates to communications networks and particularly the Internet. The invention includes an apparatus and method for publishing information to the Internet, and managing such publications to facilitate easy access by the intended recipients. The invention facilitates access to publications by enabling Internet users to subscribe to information which is published through the Internet.

The Internet comprises a network of computers interconnected by a separate communications network and using a common communications protocol. The common communications protocol and an arrangement for addressing resources or information stored or located on the network of computers provides what appears to Internet users to be seamless access to information published on the Internet virtually anywhere in the world. The interface to the Internet provided by the Worldwide Web has facilitated the development of easy-to-use browser applications which simplify Internet communications and has made the Internet available to anyone regardless of their technical knowledge.

The term "Internet" has come to be used to identify a particular network and the technology utilized by that network. However, unless otherwise noted, the term "Internet" is used in the remainder of this disclosure and the appended claims to describe the technology and networking topology employed by the particular network commonly referred to as the "Internet." Thus, the invention also has application to an "intranet", an instance of a local area network utilizing Internet technology, but which is not normally accessible by users other than those having access to the local area network.

The Internet and Worldwide Web (web) open up literally a world of information to anyone with a personal computer. Under the Worldwide Web interface, publishers or providers of information use a particular coding, Hypertext Mark-Up Language (HTML), for formatting information to be published to the web. The HTML based document or documents are stored in storage associated with an Internet or web server program operating on a computer having a constant connection to a communications system. These stored documents or files are accessible through a network wide addressing system using Universal Resource Locators (URLs) which identify a particular storage location or a particular related series of storage locations associated with a web server. URLs enable Internet users to download information which may be stored under any web server, regardless of where the web server is physically located.

Once information is published to an Internet location, it may be accessed by any Internet user having authorization to access the particular location. The user simply submits a URL, and the URL causes the identified web server, or host server to retrieve the information identified by the URL and transfer the information to the computer through which the URL was submitted. The information is transferred in a series of frames making up a web page which is displayed at the computer through which the URL was submitted. The information identifiable by a URL may actually comprise files stored at many different physical storage locations. The URL builds or defines an instruction to collect all of the related files and transfer the files to build a web page.

Commonly, users retrieve information using a graphical interface of a web browser program such as the NETSCAPE, NAVIGATOR and MICROSOFT EXPLORER web browser application programs. The HTML interface allows a document or web page identified under one Internet address or URL to include graphic links to other documents identified by other addresses. Once a first web page is loaded onto the user's computer, the user simply clicks on the graphic link to retrieve information stored at a URL address associated with the graphic link. Clicking on the graphic link causes the computer to submit the URL associated with the graphic link. This process of moving from one Internet published document or web page to the next using links embedded in the documents is commonly referred to as surfing the web.

The Internet and the Worldwide Web interface to the Internet has proven to be such a robust and convenient form of electronic communication that organizations have begun using the same technology for internal organization communications. A network utilizing Internet technology, but which is internal to an organization, such as a corporation or a governmental organization, is commonly referred to as an intranet.

Under Internet communications, published information must be located in some fashion by the user and then retrieved or down loaded to the user's location. The common method of locating Internet published information by surfing from one location to the next until a desired location is reached is time consuming and tedious. Although the Internet is highly successful at generally enabling information to be communicated, the system is very poor at organizing the information in a coherent fashion. The Internet includes no mechanism for alerting a user when important information is published or even simply facilitating access to information when it is published.

Another problem involves publishing information to the Internet through the Worldwide Web interface. Information which may need to be published may be created using a variety of software products including word processors, spread sheets, and graphics programs. In each case, the format of a particular document must be converted to a web viewable or readable format. That is, a format in which the document may be displayed at a user's terminal. Manually converting documents for Internet publication has heretofore been time consuming and has required specialized programs and knowledge.

The problems with the Internet do not end with the problems associated with publishing to the Internet and retrieving information from the Internet, that is, getting documents to the system, and then to the intended recipients. Internet communications pose significant information management difficulties. For example, prior to the present invention, Internet publishers had to maintain constant vigilance to ensure that publications are kept current and outdated publications are removed.

**SUMMARY OF THE INVENTION**

It is an object of the invention to provide a method and apparatus to overcome the above-described problems and other problems and deficiencies associated with Internet communications. More particularly, it is an object of the invention to provide a publish and subscribe system for publishing information to the Internet and enabling Internet users to subscribe to the published information, thereby improving access to the published information.

To accomplish these objects, the apparatus according to the invention includes a repository arrangement or means

and an information organizing arrangement or means. Both the repository means and the information organizing means are implemented through software operating on a computer. The computer or computers implementing the repository means and information organizing means are each connected for Internet communications through a suitable web server.

The repository means receives information elements in the form of web viewable or readable files to be published through Internet technology. Each information element received by the repository means is associated with a first network addressable location, that is, a location identified by a network address such as a URL. The repository means operates to store the information element in storage under a second network addressable location identified by a network address. As used herein, a network addressable location is simply a location which may be identified by a network address such as a URL under which information may be stored and by which the stored information may be retrieved.

Handle production means associated with the repository means produces a repository handle for each information element received by the repository means. The repository handle includes identifying information relating to the information element such as a description of the content of the information element, and the network address identifying the second network addressable location. Once the handle production means produces the repository handle, it then transfers the repository handle to the information organizing means.

The information organizing means stores the repository handle under the first network addressable location with which the information element was associated when directed to the repository means. Storing each information element at one network addressable location and then an identifier for that network addressable location under a separate network addressable location allows information elements to be conveniently organized for monitoring and retrieval by interested Internet users.

Each first network addressable location includes a listing of the handle for each information element which was associated with the particular first location when the information element was published. The handles are presented in a web page which is viewable by an authorized user using any web browser running on a remote computer having Internet access. The content descriptive identifying information for each handle is displayed in the web page in some fashion along with a graphic link which will invoke the network address for the second network addressable location where the information element is stored. Each first network addressable location may also include a list of handles from another first network addressable location to which the user has subscribed as described below. In either case, the handles appear on the web page under one or more subject divisions. For example, a first network addressable location may be set up to include handles for company reports and include as different topics, sales reports, marketing reports, and product development reports. Each topic within a first network addressable location may be further separated or divided into subject categories. For example, the sales reports topic may include categories for different regions and the product development reports topic may include categories for different development projects.

According to the invention, the information organizing means allows an Internet user to set up a first network addressable location under one URL and then subscribe to information topics set up under a different first network

addressable location under a different URL. Subscribing to a topic causes the handles of all information elements under the subscribed topic to be stored under the subscribing user's first network addressable location. Also, when a new information element is published and is associated with one first network addressable location under a subscribed topic, the information organizing means not only stores the repository handle for the published report under the associated first network addressable location, but also under each subscribing first network addressable location. To view the newly published report, the subscribing user does not have to know ahead of time that the report was published or the network address to which it was published. Rather, the subscribing user simply downloads their own first network addressable location and an identifier and link for the published report will appear there.

The invention also includes a publication arrangement or means associated with the organizing means. The publication means facilitates the publication of documents through various first network addressable locations. In the preferred form of the invention, the publication means includes processes allowing publication through non-Internet communications and through Internet communications. For publication through Internet communications, the publishing means includes programming for receiving a publish request, a file to be published representing an information element, and an association to a particular first network addressable location, all received from a remote computer through Internet communications. All of the information required for publication may be transmitted through any web browser adapted to transfer files. The publication means also preferably includes document conversion means for converting documents or files in non-web viewable formats to web viewable formats.

The preferred form of the invention further includes a reaction arrangement or means associated with the organizing means. The reaction means operates to produce a pre-defined action in response to the occurrence of an event which affects a first network addressable location. For example, a user may wish to receive an e-mail message or a page when a new quarterly sales report is published. The reaction means allows the user, from their web browser, to define the action, that is, the e-mail message or page they are to receive, and to define an event type and filter conditions which identify which report is to trigger the action. In this example, when the triggering report is published and a handle for it appears under the first network addressable location, that event satisfies the event type and filter conditions set by the user and causes the reaction means to request the predefined action, that is request the e-mail or page. The system is adapted to interface with a suitable process for performing each supported action.

The apparatus and method according to the invention facilitate the publication of information through an Internet technology network and allows others to view that information using any web browser. Publication through the repository handle listings under a first network addressable location combined with the subscription arrangement included in the preferred form of the invention allows users to locate topics and automatically obtain links to information elements under such topics. This is in contrast to prior Internet publishing where a user had to surf haphazardly from one location to the next to locate published information. Furthermore, the reaction means associated with the invention allows users to obtain notification when information is published under a particular topic, unlike prior Internet publishing which included no reliable means by which an

interested user could obtain notification when useful information was available.

These and other objects, advantages, and features of the invention will be apparent from the following description of the preferred embodiments, considered along with the accompanying drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagrammatic representation of a network publish and subscribe system embodying the principles of the invention.

FIG. 1A is a diagrammatic representation of the transfer of information to the repository means and information organizing means according to the invention.

FIG. 2A is a representation of a primary display for a first network addressable location embodying the principles of the invention.

FIG. 2B is a representation of an alternative primary display for a first network addressable location.

FIG. 3 is a block diagram illustrating the elements of a publish and subscribe system according to the invention.

FIG. 4 is a flow chart showing the manner in which requests submitted by Internet users are processed according to the invention.

FIG. 5 is a diagrammatic representation of a display according to the invention through which an Internet user may publish an information element.

FIG. 6 is a flow chart showing the process steps performed for publishing an information element.

FIG. 7A is a diagrammatic representation of a display an Internet user may use according to the invention to subscribe to information elements published through a particular first network addressable location.

FIG. 7B is a diagrammatic representation of an alternative display through which an Internet user may subscribe to information elements published through a particular first network addressable location.

FIG. 8 is a flow chart illustrating the process steps performed in response to a subscription request.

FIG. 9A is a diagrammatic representation of a display through which an Internet user may submit a reaction request according to the invention.

FIG. 9B is a diagrammatic representation of a display for defining reaction attributes.

#### DESCRIPTION OF THE PREFERRED EMBODIMENTS

FIG. 1 illustrates the basic components and interactions of a network publish and subscribe system 12 embodying the principles of the invention. The system 12 is comprised of various software components or processes operating on a digital computer. The core components of the invention comprises an information organizing means 14 and a repository means 15 both programmed on computer or system server 16. Although the information organizing means 14 and repository means 15 are illustrated in FIG. 1 as operating on a single computer 16, those skilled in the art will readily appreciate that the processes may be operating within a suitable distributed computing environment in which the processes are actually being executed by more than one computer or processor. Throughout this disclosure, a process operating within a distributed computing environment should be considered equivalent to such a process operating on a single computer, such as computer 16.

Referring to FIGS. 1 and 1A, the repository means 15 operates to receive an information element 18 to be published along with identifying information 19 for the information element, and then store the information element under a network addressable location identified in some fashion, such as by URL<sub>2</sub>, for example. The information element 18 is preferably stored in a mass storage device 20 associated with computer 16.

The computer or system server 16 through which the repository means 15 is operating is connected for communications through a communications network and particularly, an Internet communications network. The term "network addressable location" as used in this disclosure and the appended claims means a location which is addressable or identifiable by an addressing arrangement employed by the particular communications network. Referring particularly to Internet technology, the addressing arrangement may comprise the URL system. Thus, the network addressable location refers to a "location" from which information including files or documents may be retrieved using a network address. It will be appreciated that the "location" is not necessarily a single, contiguous area of memory in mass storage. Rather, an information element 18 may be built from many separate files or documents filed in separate mass storage locations and recalled by a particular URL or other network address. In the Internet example, the URL serves as a command to retrieve the separate components of an information element 18 and assemble the components to create the information element.

The invention will be disclosed herein as implemented utilizing Internet communications technology. However, those skilled in the art will appreciate that the invention may be implemented using other communications technology which provides functionality similar to that provided by the Internet.

The repository means 15 receives the information element 18 and the identifying information 19 from a publication means or arrangement which will be described below. The information element 18 may originate from a network computer 22 communicating with the system server 16 by some means or from a remote computer 23 through Internet communications. When the information element 18 is received by the repository means 15, the information element has already been associated with another network addressable location, URL<sub>1</sub> in FIG. 1A for example, different from where the information element is stored by the repository means. This separate network addressable location (e.g., URL<sub>1</sub>) associated with the information element 18 by the publisher will be referred to in this disclosure and the accompanying claims as a first network addressable location 30. The network addressable location (e.g., URL<sub>2</sub>) at which the information element 18 is stored by the repository means 15 will be referred to as a second network addressable location 31.

The identifying information 19 received by the repository means 15 is also defined or produced through the publication means. The identifying information 19 preferably includes a title of the information element 18, a description of the information element, a creation time, and a source of the information element. Some information elements may relate to an event such as an appointment. In such cases, the identifying information 19 for the information element 18 may include start and end times for the event.

Upon receipt and storage of the information element 18 by the repository means 15, a handle production means or arrangement 24 associated with the repository means pro-

duces a repository handle 26 for the information element 18. The repository handle 26 includes the identifying information 19 and the network address (e.g., URL<sub>2</sub> or elements from which URL<sub>2</sub> may be developed when required) for the second network addressable location 31 within which the information element 18 has been stored. Thus, the repository handle 26 includes all information required to identify the information element 18 and to recall or retrieve the information element through the network addressing arrangement.

After the handle 26 is produced for an information element 18 being published, the handle the production means 24 transfers the repository handle to the information organizing means 14. The information organizing means 14 stores the repository handle 26 under the first network addressable location 30 (e.g., URL<sub>1</sub>) with which the particular information element 18 was associated when published. As with the storage performed by the repository means 15, the information organizing means 14 preferably uses mass storage 27 associated with the computer 16 through which the organizing means is operating. Although shown as separate storage devices, those skilled in the art will readily appreciate that the mass storage devices 20 and 27 used by the repository means 15 and information organizing means 14 may be the same storage device.

The repository means 15 stores the actual information element 18 published to the system 12. The information element 18 may be a document such as a report or letter, or a file such as a database file, or any other type of information. Many different information elements 18 may be associated with a single first network addressable location 30 such as URL<sub>1</sub> in FIG. 1A. Thus, the handles 26 for the different information elements 18 are stored by the information organizing means 14 under the common first network addressable location. When the first network addressable location 30 (e.g., URL<sub>1</sub>) is retrieved, the handle 26 for each associated information element 18 is also retrieved. Since each handle 26 includes all information necessary to describe and identify a particular associated information element 18 and to retrieve the information element itself, a user may use the information stored under the first network addressable location 30 to locate and retrieve information elements of interest.

The information organizing means 14 is capable of storing handles 26 under any number of different first network addressable locations 30. Each such network location represents a location through which related information elements or information elements of interest may be accessed. The address (e.g., URL<sub>1</sub>) for retrieving handles 26 stored under one of the first network addressable locations also preferably retrieves information which produces a display in a particular format and including controls for performing various functions. These elements together, the handles 26 stored under a first network addressable location 30 and the display information retrieved along with the handles, build a primary display or interface comprising a web page of web viewable information. The primary display will be discussed in detail referring to FIGS. 2A and 2B.

As shown in FIG. 1, the publish and subscribe system 12 preferably includes a plurality of system servers 16 all interconnected using Internet technology. A web server process 32 is thus associated with each system server 16 and each remote computer 23. A user at a remote computer 23 may use the Internet communications protocol, to retrieve and view all handles 26 stored under any first network addressable location 30 stored under any system server 16. The user retrieves a particular first network addressable

location 30 by submitting a network address (e.g. URL<sub>1</sub>) for the desired first network addressable location.

The URL may be submitted conveniently through web browser 33. Using a web browser 33, the contents first network addressable location displays as a web page under the remote user's web browser application. The web page display lists the viewable portions of the handle 26 associated with each information element 18 associated with the particular first network addressable location 30 and also, includes the network address (e.g. URL<sub>2</sub> in FIG. 1A) for the second network addressable location 31 under which a particular information element 18 is stored. Preferably, the network address for the information element 18 is included under a graphic link on the display. From the web page for the first network addressable location 30, the user may view an information element 18 associated with a handle 26 displayed on the page by clicking on the graphic link. This action invokes the network address of the second network addressable location 31 (e.g. URL<sub>2</sub>) causing the information element 18, to be retrieved to the user's remote computer 23 where it may be viewed if it is in a web viewable format.

FIG. 1 illustrates the preferred implementation of the invention in which the network communications technology comprises Internet technology. This communications technology allows any number of users to access information published through the first network addressable locations 30. Using the common addressing system provided by the URL system, the invention allows users to view locations 30 and information elements 18 which are stored anywhere on the system, that is, on any system server 16. This accessibility combined with the subscription and reaction processes described below allows the system 12 according to the invention to greatly increase accessibility to information.

FIGS. 2A and 2B illustrate two alternative primary displays or graphical user interfaces (GUIs) for a first network addressable location (30 in FIG. 1). A display may appear under any web browser through which the first network addressable location 30 may be retrieved. The primary display in FIG. 2A illustrates a framed display, while FIG. 2B illustrates an unframed display. Referring to FIG. 2A, the display includes a toolbar 40, an index 41, and a content area 42. Although not shown in the figures, the primary display may also include a header and a footer.

Referring still to FIG. 2A, the illustrated preferred toolbar 40 includes a topic selector 43 and a plurality of operation buttons 44. The topic selector 43 is used to select a topic division under which handles 26 may be grouped. The operation buttons 44 may include display options allowing the display to include, for example, all handles 26 or just those relating to a certain period of time, such as a particular day, week, or month. Other operation buttons 44 may facilitate publishing to the system, set-up operations such as first network addressable location, display, topic, subscription, and permission set-up, a search operation, a refresh operation, and a help operation.

Each handle 26 is displayed in the index 41. The current date may be displayed at the top of the index 41 portion of the display. Each handle 26 includes the identifying information portion of the handle, and preferably a graphic link which will invoke the address or URL for the second addressable location where the associated information element is stored. The graphic link may be the title portion of the handle 26 such that the action of clicking on the title invokes the URL for the information element associated with the particular handle. The index also includes a scroll-bar provided by the browser application to allow the user to

scroll down through the index to view the various handles 26 which are listed.

The handles 26 are grouped according to a particular subject division or topic. The primary subject divisions or topics may be broken down into further divisions or categories. Any of these subdivisions or categories may be defined so as to appear under more than one topic. The user sets up the topics and categories or other subject divisions included in the particular display. In the illustrated displays, the index includes handles under one topic at a time and topics are chosen through the topic selection button 43 and menu associated with the toolbar 40.

The display shown in FIG. 2A includes a content area 42 for displaying an information element retrieved from a second network addressable location. The content area 42 includes a scrollbar provided by the browser application to allow the user to scroll through long documents displayed to the content area 42.

The alternative display shown in FIG. 2B illustrates a non-framed display. This alternative display includes the toolbar 40 with topic selection device 43, index 41, and if desired, a header and footer (not shown), all similar to those elements shown in FIG. 2A. However, the non-framed display illustrated in FIG. 2B includes no content area. Handles 26 under a chosen topic are simply listed across the entire display. When a link within a handle is clicked to retrieve an information element, the information element appears over the entire display replacing the primary display. The user may return to the primary display by moving to the previous page function of the web browser.

The publish and subscribe system according to the invention includes a permissions system (not shown) which defines each user's permission to perform particular system operations. The permissions granted to various users are preferably stored in a data table maintained by the system server 16. For each requested operation, the system server 16 queries the permissions data table to determine whether the particular user has permission to perform the requested operation. In the preferred form of the invention, the user supplies a user name and password in order to retrieve a particular first network addressable location, and then permissions for further operations are analyzed for that user name and password.

FIG. 3 illustrates the architecture of the preferred implementation of the invention. The system architecture may be separated into the subscribe components 50, publish components 51, and system components 52. The subscribe components 50 enable a user to subscribe to information elements which may be published through the system. The publish components 51 enable a user to publish an information element to the system.

Subscribing according to the invention involves a process in which a user incorporates a subject division (topic) or even a handle (26 in FIGS. 1 and 1A) from one first network addressable location (30 in FIG. 1) into another. When a user subscribes to a topic of first network addressable location X, for example, that topic and all handles under the topic are incorporated into the subscriber's first network addressable location. That is, all handles under the subscribed topic under first network addressable location X are referenced under first network addressable location representing the subscriber's first network addressable location. Furthermore, any future information element which is associated with first network addressable location X when published, will in addition to having the associated handle stored in the first network addressable location X, will also,

have the handle stored under the subscriber's first network addressable location. Thus, a subscribing user need only view the web page display for their own first network addressable location, to obtain links to other information of interest which has been published. The subscribing user obtains these links automatically without needing to have any knowledge of the publication.

A user may subscribe to information elements published to the system through the user's web browser 33 using the primary display, potentially through other applications 60 supported by the system, and by remote system servers 61. In each case, the particular component, browser 33, application 60, or remote system server 61 submits a subscribe request to subscription means associated with the organizing means 14 (FIG. 1) operating on the system server 16. The subscription request is submitted through the web server 32 and a web server interface 64. The web server 32 passes the request from the submitter using Internet communications protocol. The web server interface 64 receives the request under the Internet protocol and transforms the request to a system request form, that is, a form acceptable to the system server 16. The transformed request is then submitted to a system server interface 65 to invoke commands or functions to be performed by subscription means programmed in the system server 16. The commands or functions invoked by the subscription request include a search for permission to subscribe as requested, the addition of information to a subscription table maintained in a database 68 by the system server 16, and referencing of the subscribed topic to the subscriber's first network addressable location. This subscription process is discussed with reference to FIGS. 7A, 7B and 8 below.

The illustrated architecture shows three avenues for publishing an information element (18 in FIG. 1) to the system under the publish components 51. The first avenue is through Internet communications technology. Under this publication avenue, the user submits a publish request through their browser 33 using a display such as the primary display shown in FIGS. 2A and 2B. The publish request is transmitted in the network protocol through the web server 32. The web server interface 64 transforms the publish request from the network protocol to a system form. The request in system form is then submitted to the system server interface 65 which invokes the commands or functions which are to be performed by publication means programmed in the system server 16. The commands for publication include obtaining the identified document or other information element to be published, storing the information element and the handles to the information element, and searching for reactions to be performed. The publication means and process is discussed in detail with reference to FIGS. 5 and 6.

The two other avenues for publication require non-Internet based communications to the system server 16. The system according to the invention may be adapted to communicate with various applications 70 or through resource gateways 71 for resources such as facsimile machines, and scanners, for example, to receive information elements to be published. An example of an application 70 which may be used to submit documents to be published and publication requests comprises the delivery system disclosed in U.S. patent application Ser. No. 08/330,390, entitled APPARATUS AND PROCESS FOR DISTRIBUTING INFORMATION USING AN ELECTRONIC PACKAGE REPRESENTATION, now U.S. Pat. No. 5,870,089, which is incorporated herein by reference. Publication through an application 70 or gateway 71 according to the invention

comprises submitting the publication request directly from the gateway or application to the system server interface 65 in system form to invoke the various commands or functions to be performed by the publication means associated with the system server 16.

The system server 16 includes the repository means 15 and organizing means 14 and their various components for performing the various functions of the system. The system server 16 also maintains the database 68 which forms a component of the organizing means 14 and which includes data tables containing linking information, first network addressable location information, permission information, and subscription information. A document management system interface 73 and document management system 74 are components of the repository means 15 and are utilized by the system server 16 for storing and retrieving documents published through the system. An output interface 76 to an output system 77 is provided for directing documents to resources such as printers and resources for performing reactions such as sending e-mail messages and pager messages. These output components preferably comprise the delivery system disclosed in U.S. patent application Ser. No. 08/330,390. Also, the system includes an interserver communications process 79 for communicating with remote servers using Internet communications.

FIG. 4 illustrates the process utilized by the invention for handling requests submitted via an Internet communication. The process applies for all Internet submitted requests whether the request originates through an Internet browser 33, some other type of application 60 adapted to submit requests via the Internet, or a remote system server 61, all shown in FIG. 3. The submitted request is first communicated to the system via the Internet communication and then at step 82 received at the web server 32 associated with the local or host system server 16 shown in FIG. 3. The web server 32 transfers the network protocol request to the web server interface 64 which, at step 83, transforms the network protocol request to a system request, that is, to a format acceptable to the system server interface 65. The web server interface 64 then directs the system request to the system server interface 65 which at step 84 communicates the commands and functions required by the request along with any required data to the system server 16 for execution at step 85.

FIG. 5 may be used to describe how a publish request is formulated at a browser 33, while FIG. 6 illustrates the process steps required for publishing an information element 18 according to the invention. The dialogue or GUI shown in FIG. 5 appears over the primary display or page when the user chooses a publish tool on the primary display toolbar 40 (FIGS. 2A and 2B). This dialogue, along with all other dialogues chosen through the toolbar 40 are actually transferred to the remote computer 23 along with data required to produce the primary display and the information contained within the first network addressable location. The dialogue shown in FIG. 5 requests information required to publish an information element (18 in FIG. 1), including identifying information for the information element to be published, the location of the information element, and attributes of the publication. The identifying information preferably includes the title of the information element to be entered at box 90 and a description of the information element to be entered at box 91. The publication attributes include a subject division or category selected at boxes 92, start and end information to be entered at boxes 93 if the information element refers to an appointment, and finally display dates at boxes 94 and a removal date at box 95. The location of the document to

be published is identified by a file name or a URL at box 96 in the dialogue.

When the user has entered the information requested in the dialogue shown in FIG. 5 and clicks the "OK" button, the primary display or interface illustrated at FIGS. 2A and 2B formulates or builds a publish request in the form of a specialized URL. This publish request is processed like any other Internet submitted request as shown in FIG. 4.

The operations or steps performed at step 85 in FIG. 4 for a publication request are set out in FIG. 6. At step 100, the publication means programmed in server 16 in FIG. 3 first extracts from the system request all of the information required to publish the particular document, including the identifying information, the information element to be published (or the URL to the information element), and publication attributes, including the first network addressable location under which the handle for the information element is to be stored. It will be noted that the dialogue shown in FIG. 5 does not request information on the first network addressable location where the handle is to be stored. Rather, the handle is stored under the first network addressable location associated with the primary display from which the publish dialogue was chosen.

After extracting all of the required information from the system request, the process preferably includes at step 101 converting any file representing the information element to a web viewable format, in the event that the document is not already in such a format. This conversion may be performed in any suitable manner. Once any document conversion is performed, the information element and the identifying information are submitted to the repository means (15 in FIG. 1) at step 102 and the repository means then at step 103 stores the information element at an available second network addressable location. The second network addressable location may be defined by the repository means 15. Once the second network addressable location is defined, the handle production means (24 in FIGS. 1 and 1A) then creates a repository handle (26 in FIG. 1A) for the information element.

Once the handle is created, the organizing means at step 105 stores the handle associated with the information element under the first network addressable location specified by the publisher. Subscription means at step 105 then updates a pending subscriptions data table in the database 68 to enable subscriber locations to recognize that the published information element has been published and a subscription data table must be updated. The subscription means comprises a program associated with the organizing means, for causing the system server 16 to perform the data table updates. When the subscriber loads their first network addressable location, a marker in the pending subscription data table causes the handle of the published element to be added at the subscriber location.

The final steps of the publication process involve reactions which will be discussed in detail below with reference to FIGS. 9A, 9B, and FIG. 10. At step 107, reaction means associated with the information organizing means (14 in FIGS. 1 and 1A) searches a reaction data table maintained in the system server database 68 to determine if the publication matches any triggering event type and satisfies an associated filter stored in the reaction table. If all of the reaction conditions are satisfied, the reaction means at step 108 produces a request to perform the predefined action. If the action must be performed by a resource external to the system 12, the system server 16 sends a command through the output interface 76 to the particular output resource 77 (FIG. 3) for performing the action.

The processes for submitting a subscription request through Internet communications may be described with particular reference to FIGS. 7A, 7B, and 8. FIGS. 7A and 7B show alternative displays or GUIs associated with the primary display shown in FIGS. 2A and 2B. The dialogue shown in FIG. 7A is invoked when a subscriber clicks a subscribe tool associated with a display showing their own first network addressable location. Since the subscriber's first network addressable location is known, the only information required for the subscription is an identification for the publisher's first network addressable location and the subject division (topic) under the publisher's location to which a subscription is requested. This information may be added at boxes 120 and 121. The dialogue shown in FIG. 7B is invoked from a system subscribe tool which may appear at the bottom of the index in the primary display for a publisher's first network addressable location. In this case, the identity of the publisher's location is known and all that is required is that the subscriber select the publisher's topic to which they desire a subscription at box 122 and then identify the subscriber's own first network addressable location at 123. Clicking the "OK" button from either dialogue 7A or 7B causes the primary display to formulate or build a subscribe request. This subscribe request is submitted according to the steps set out in FIG. 4.

Referring to FIG. 8, the subscription operation performed by subscription means programmed in the system server 16 includes at step 130 extracting the relevant data from the incoming request. For subscriptions, the data includes the first network addressable location being subscribed to, along with the subject division (in the illustrated case a topic division), and the subscriber first network addressable location where the information elements and topic being subscribed to are to appear. The subscription means then at step 131 adds the data to the subscription table maintained in the system server database 68. Finally, at step 132 in FIG. 8, the subscription means copies any handles already existing under the subscribed topic under the first network addressable location to which the subscription applies and stores those handles and subject division under the subscriber first network addressable location.

FIGS. 9A and 9B illustrate primary display dialogues which may be used to set up a reaction according to the invention. Reaction means comprising programming operating on the system server (16 in FIG. 3) operates to detect a specified or predefined triggering event satisfying any predefined conditions and then requests or invokes some predefined action in response. The types of actions which may be predefined are limited by the output resources available to the system server 16, e-mail message, pager messages, faxes, or other actions may be available to be set up by a system administrator. The dialogue or GUI shown in FIG. 9A is invoked preferably through a "set-up reaction" item on a set-up menu available under a set-up tool on the primary display toolbar 40 (FIGS. 2A and 2B). The dialogue requests a reaction name or identifier at box 140 which may be used to identify the particular reaction, an event type or triggering event type at box 141, and an action type at box 142. Once this information is entered and the "next" button is activated, a dialogue such as that shown in FIG. 9B appears over the primary display. This dialogue allows a subject division or category to be specified at box 143 so that only events effecting that subject division will invoke the desired action. The information from the event options portion of the dialogue in FIG. 9B is used by the reaction means to set up a filter which must be satisfied before the predefined action is requested. The dialogue shown in FIG.

9B also requests information to define the particular action. This particular dialogue is specifically for an e-mail action and requests the mail message to be included in the e-mail, the subject, and the recipient for the e-mail at boxes 145, 146, and 147, respectively.

Once the information is added and the user clicks the "OK" button on the dialogue shown in FIG. 9B, the primary display creates a subscription request in the form of a specialized URL and then transmits that request URL to the system server 16 via Internet communications. This request is handled similarly to any Internet submitted request as shown in FIG. 4. On receipt of the subscription request at the system server 16, the system server extracts the information for the particular reaction and stores the information under the reactions data table in the system server database 68. The reaction means searches the reaction table as illustrated at step 107 in FIG. 6 each time a potential triggering event occurs. If the event type matches a triggering event type stored in the reaction table, and if any filter conditions specified in the table are satisfied, then the reaction means produces a request directing the specified output device to perform the predefined action as illustrated at step 108 in FIG. 6.

The preferred form of the invention also includes an arrangement or means for notifying a user of new handles which are stored under a first network addressable location which the user has accessed. This notification means is activated each time a user accesses or loads the contents of a first network addressable location. When the location is loaded to the user's computer, the notification means produces a notification request and communicates the request, preferably in the form of a URL, to the system server from which the location was loaded. This notification request causes the server 16 to produce a query to search for handles whose last update time is more recent than the last time that the particular user accessed the particular first network addressable location. The system server 16 will then return to the user location instructions to change the display (FIGS. 2A and 2B) in some desired fashion to indicate which handles contain new information. For example, the display may be changed to highlight the handles associated with new information or provide some image next to such handles.

In the preferred form of the invention, interactions between a user at a remote computer 23 and system server 16, as well as interactions between system servers are based upon an attribute-value paradigm. Under this paradigm, all data that is communicated from the remote user to the system server 16 and back again is done as a set of attributes that have one or more associated values for each. The system employs an attribute value pair or AVP library comprising predefined attribute names which may be associated with zero or more values. AVPs are formed into sets or lists which are passed between the communicating entities to communicate requests and data.

As shown in FIG. 1, many system servers 16 may be included in a system 12 embodying the principles of the invention. Users generally set up their first network addressable locations at a particular system server identified as their host or local system server. Each local system server is responsible for updating local first network addressable locations and keeping subscriptions relating to these first locations up-to-date. When a local user subscribes to a first network addressable location which is local to another system server or a remote system server, the subscription means associated with the local server causes the local server to contact the remote server to update the subscription when necessary.

In the preferred form of the system servers 16 cache information from remote system servers as necessary to speed system processes. For example, if a user has a first network addressable location at their local system server A, and enters a subscription for a first network addressable location local to remote system server B, the subscription means associated with system server A will contact system server B to update the subscription data tables maintained at system server B. System server A will also cache the information from system server B so that repeated access to the information will be performed quickly.

The system also preferably includes a web page based administrator's display or interface through which a particular first network addressable location may be defined. Using a set-up dialogue associated with the administrator's display, a user having special authorization to set-up first network addressable locations may enter the required information to define the attributes of the first location to be created. The attributes may include an identifier or name for the location, and permissions information such as whether the location is private and accessible only by the user, or public and accessible by anyone having access to the work. Once the required information is entered, the display formulates a location set-up request which is processed like any other system request as discussed above with reference to FIG. 4.

The above described preferred embodiments are intended to illustrate the principles of the invention, but not to limit the scope of the invention. Various other embodiments and modifications to these preferred embodiments may be made by those skilled in the art without departing from the scope of the following claims.

We claim:

1. An apparatus for publishing an information element through a first network addressable location of a communications network and improving access to such published information from a terminal having access to the network, the apparatus comprising:

(a) repository means programmed in a computer for receiving the information element and storing the information element at a second network addressable location, the computer being connected to the communications network;

(b) handle production means programmed in the computer for receiving identifying information relating to the information element and for producing a repository handle for the information element stored by the repository means, the repository handle including identifying information relating to the information element and a network address for the second network addressable location; and

(c) information organizing means programmed in the computer for retrieving from the handle production means the repository handle associated with the information element within the second network addressable location and storing said repository handle under the first network addressable location.

2. The apparatus of claim 1 wherein the information organizing means includes:

(a) subscription means programmed in the computer, the subscription means for receiving over the communications network a subscription request identifying characteristics of a future repository handle to be stored under the first network addressable location and, in response to the production of the future repository handle, copying the future repository handle and then storing said copied repository handle under a subscriber network addressable location.

3. The apparatus of claim 2 wherein the communication network comprises the Internet and wherein the information organizing means has associated therewith:

(a) an incoming web interface for receiving Internet communications from a remote computer.

4. The apparatus of claim 3 wherein:

(a) the first network addressable location includes a first topic division within which the future repository handle is to be associated with a plurality of other repository handles which are each associated with a different information element; and

(b) the subscription request identifies the characteristics of the future repository handle stored under the first network addressable location by identifying the first topic division.

5. The apparatus of claim 4 wherein the first network addressable location includes a plurality of other topic divisions in addition to the first topic division.

6. The apparatus of claim 4 wherein:

(a) the first topic division includes a first category division within which the future repository handle is to be associated with a plurality of additional repository handles which are each associated with an additional information element; and

(b) the subscription request identifies the characteristics of the future repository handle to be stored under the first network addressable location by identifying the first category division.

7. The apparatus of claim 6 wherein the first topic division includes a plurality of other category divisions in addition to the first category division.

8. The apparatus of claim 3 wherein the information organizing means further includes:

(a) publication means programmed in the computer, the publication means for receiving from the remote computer via Internet communications a file containing the information element to be published and in addition to said file a publish request, and for extracting from the publish request the identifying information relating to the information element and a definition for the first network addressable location, and for passing said identifying information to the handle production means.

9. The apparatus of claim 8 wherein the publication means is also for extracting display information from the publish request, the display information associating the handle for the information element with a subject division under the first network addressable location.

10. The apparatus of claim 8 further comprising:

(a) format conversion means associated with the publication means for identifying the format of the file containing the information element and converting the file to a web viewable format.

11. The apparatus of claim 3 further comprising:

(a) publication means programmed in the computer, the publication means for receiving from the remote computer via Internet communications a URL associated with the information element to be published and a publish request, and for extracting from the publish request the identifying information relating to the information element and a definition for the first network addressable location, and for passing the identifying information to the handle production means.

12. The apparatus of claim 1 further comprising:

(a) publication means programmed in a network computer in communication with the computer, the publication

17

means (i) for receiving a publish request including a file containing the information element to be published, and for (ii) extracting from the publish request the identifying information relating to the information element, the file containing the information element to be published, and a definition for the first network addressable location, and for (iii) passing said identifying information to the handle production means.

13. The apparatus of claim 1 wherein the information organizing means further includes:

(a) reaction means programmed in the computer, the reaction means for producing an action request for a predefined action in response to the occurrence of an event affecting the first network addressable location.

14. The apparatus of claim 13 wherein the reaction means includes:

(a) reaction storage means for storing in a reaction data table a reaction identifier, a triggering event type, a filter, and an action definition defining the predefined action;

(b) search means for searching the reaction data table in response to a potential triggering event to determine if the potential triggering event matches the triggering event type and satisfies the filter stored in the reaction data table; and

(c) action request means for producing the action request if the potential triggering event matches the triggering event type and satisfies the filter.

15. The apparatus of claim 14 wherein the reaction means further includes:

(a) reaction definition means for receiving from a remote computer via Internet communication a reaction set-up request and for extracting from the reaction set-up request the reaction identifier, the triggering event type, the filter, and the action definition.

16. A method for publishing an information element so as to be accessible from a first Internet location, the method comprising the steps of:

(a) receiving a publish request for the information element, the publish request including the information element, identifying information relating to the information element, and a URL for the first Internet location;

(b) producing a repository handle for the information element, the repository handle including a URL for the second Internet location and the identifying information relating to the information element;

(c) storing the repository handle for the information element under the first Internet location; and

(d) in response to the receipt of the publish request, storing the information element at a second Internet location.

17. A process for publishing information to a network addressable location of a communications network and improving access to such published information from a terminal having access to the network, the method comprising the steps of:

(a) associating an information element to be published with a first network addressable location;

(b) transferring the information element to a repository for storage therein, the repository comprising a computer storage device associated with a computer which is connected to the communications network;

(c) storing the information element at a second network addressable location at the repository;

18

(d) in response to the transfer of the information element to the repository, producing a repository handle for the information element, the repository handle including a network address for the second network addressable location at which the information element is stored, and also including identifying information relating to the information element; and

(e) storing the repository handle under the first network addressable location.

18. The process of claim 17 wherein the communications network comprises the Internet communications network.

19. The process of claim 18 further comprising the steps of:

(a) receiving via Internet communications a subscription request identifying characteristics of a future repository handle to be stored under the first network addressable location; and

(b) in response to the production of the future repository handle, copying the future repository handle and then storing said copied repository handle under a subscriber network addressable location.

20. The process of claim 19 further comprising the steps of:

(a) in response to the subscription request, storing the characteristics of the future repository handle and the subscriber network addressable location in a subscription data table; and

(b) searching the subscription data table in response to the production of the repository handle to determine if the repository handle comprises the future repository handle.

21. The process of claim 20 wherein the characteristics of the future repository handle to be stored in the subscription data table comprise a subject division under the first network addressable location.

22. The process of claim 17 further comprising the steps of:

(a) receiving a publish request at a computer associated with the repository; and

(b) extracting from the publish request the identifying information relating to the information element and a definition for the first network addressable location.

23. The process of claim 22:

(a) wherein the publish request is received via Internet communication.

24. The process of claim 22 further comprising the step of:

(a) extracting from the publish request display information associating the repository handle for the information element with a subject division under the first network addressable location.

25. The process of claim 17 further comprising the step of:

(a) identifying the format in which the information element is transferred to the repository and converting the format to a web viewable format.

26. The process of claim 18 further comprising the step of:

(a) producing an action request for a predefined action in response to the occurrence of an event effecting the first network addressable location.

27. The process of claim 26 further comprising the steps of:

(a) receiving from a remote computer via Internet communication, a reaction set-up request;

(b) extracting from the reaction set-up request a reaction identifier, a triggering event type, a filter, and an action definition;

19

- (c) storing in a reaction data table the reaction identifier, the triggering event type, the filter, and the action definition;
- (d) in response to the occurrence of a potential triggering event, searching the reaction data table to determine if the potential triggering event matches the triggering event type and satisfies the filter; and
- (e) producing the action request if the potential triggering event matches the triggering event type and satisfies the filter.

28. An apparatus for publishing information to an Internet location and improving access to such published information by an authorized Internet user:

- (a) a repository computer system including a digital information storage device, the repository computer system adapted to receive identifying information relating to an information element and the information element which is to be accessible through a first Internet location and to store the information element at a second Internet location, the repository computer system including an Internet server process and operating to produce a repository handle when the information element is stored in the storage device, the repository handle including a URL for the second Internet location and identifying information relating to the information element; and
- (b) a handle organizing computer system including a digital information storage device and an Internet server process, the handle organizing computer system adapted to receive from the repository computer system the repository handle associated with the information element and to store the repository handle under the first Internet location.

29. A program product for operating on a computer having access to a communications network, the program product for publishing an information element through a first network addressable location of the communications network and improving access to such published information from a terminal having access to the network, the program product comprising:

- (a) a computer-readable storage medium;
- (b) repository means stored on the medium for enabling the computer to receive the information element and store the information element at a second network addressable location;
- (c) handle production means stored on the medium for enabling the computer to receive identifying information relating to the information element and for producing a repository handle for the information element, the repository handle including identifying information relating to the information element and a network address for the second network addressable location; and

20

- (c) information organizing means stored on the medium for enabling the computer to obtain from the handle production means the repository handle associated with the information element stored at the second network addressable location and store said repository handle under the first network addressable location.

30. The program product of claim 29 wherein the information organizing means includes:

- (a) subscription means stored on the medium, the subscription means for enabling the computer to receive over the communications network a subscription request identifying characteristics of a future repository handle to be stored under the first network addressable location and, in response to the production of the future repository handle, enabling the computer to copy the future repository handle and then store said copied repository handle under a subscriber network addressable location.

31. The program product of claim 30 wherein:

- (a) the first network addressable location includes a first topic division within which the future repository handle is to be associated with a plurality of other repository handles which are each associated with a different information element; and
- (b) the subscription request identifies the characteristics of the future repository handle stored under the first network addressable location by identifying the first topic division.

32. The program product of claim 29 wherein the information organizing means further includes:

- (a) publication means stored on the medium, the publication means for enabling the computer to receive from a remote computer via Internet communications a file containing the information element to be published and in addition to said file a publish request, and to extract from the publish request a definition for the first network addressable location and the identifying information relating to the information element, and to pass said identifying information for use by the handle production means.

33. The program product of claim 29 wherein the information organizing means further includes:

- (a) reaction means stored on the medium, the reaction means for causing the computer to produce an action request, the action request directing the computer to invoke a predefined action in response to the occurrence of an event affecting the first network addressable location.

\* \* \* \* \*



US006167408A

**United States Patent [19]****Cannon et al.****[11] Patent Number:** 6,167,408**[45] Date of Patent:** Dec. 26, 2000

**[54] COMPARATIVE UPDATES TRACKING TO SYNCHRONIZE LOCAL OPERATING PARAMETERS WITH CENTRALLY MAINTAINED REFERENCE PARAMETERS IN A MULTIPROCESSING SYSTEM**

**[75] Inventors:** David Maxwell Cannon; David Romney Crockett, both of Tucson, Ariz.

**[73] Assignee:** International Business Machines Corporation, Armonk, N.Y.

**[21] Appl. No.:** 09/144,425

**[22] Filed:** Aug. 31, 1998

**[51] Int. Cl.<sup>7</sup>** G06F 12/00

**[52] U.S. Cl.** 707/203; 707/200; 707/201; 707/202

**[58] Field of Search** 707/9, 200-204; 709/201, 220, 221, 225, 229, 248; 713/200-202; 395/712; 712/32

**[56] References Cited**

**U.S. PATENT DOCUMENTS**

5,475,819	12/1995	Miller et al.	709/203
5,537,642	7/1996	Glowny et al.	713/200
5,546,595	8/1996	Norman et al.	710/10
5,606,693	2/1997	Nilsen et al.	707/10
5,652,908	7/1997	Douglas et al.	714/4
5,668,986	9/1997	Nilsen et al.	707/10
5,956,719	9/1999	Kudo et al.	707/10

**OTHER PUBLICATIONS**

Orfali et al., "essential Client/Server Survival Guide", 1994, pp. 430-474.

*Primary Examiner*—John A. Follansbee

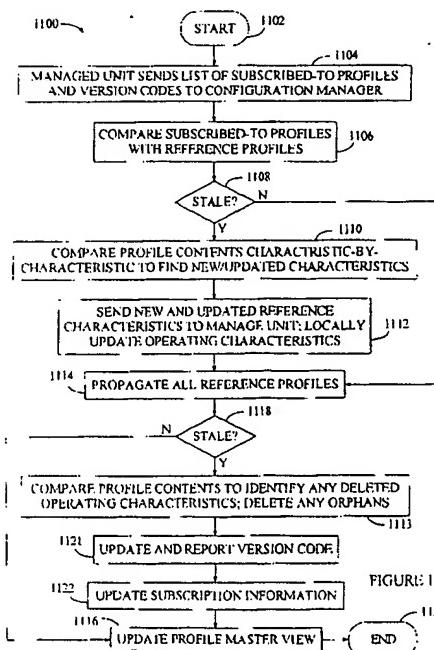
*Assistant Examiner*—Stacy Whitmore

*Attorney, Agent, or Firm*—Dan Hubert & Assoc.

**[57] ABSTRACT**

In a multiprocessor system, a configuration manager maintains various reference parameters that are selectively copied by subordinate managed units to form local operating parameters, which subsequently govern operation of these managed units. A comparative technique is employed to track reference parameter updates, and synchronize each local operating parameter counterpart accordingly. At the configuration manager, reference parameters include reference profiles and reference characteristics. Each reference profile specifies one or more of the reference characteristics. At each managed unit, the operating parameters include subscribed-to profiles and operating characteristics; both are initially copied from the configuration manager's reference profiles/characteristics. Each local operating profile specifies one or more of the operating characteristics. Each managed unit operates according to its locally maintained operating characteristics. When certain update criteria are satisfied, a managed unit and the configuration manager cooperatively synchronize the managed unit's local operating profiles and characteristics with the configuration manager's reference profiles and characteristics. This involves comparing the reference and operating profiles to identify new, updated, or deleted operating characteristics. Also, the local operating profiles and operating characteristics may be cross-referenced to identify any "orphan" characteristics for deletion.

25 Claims, 8 Drawing Sheets



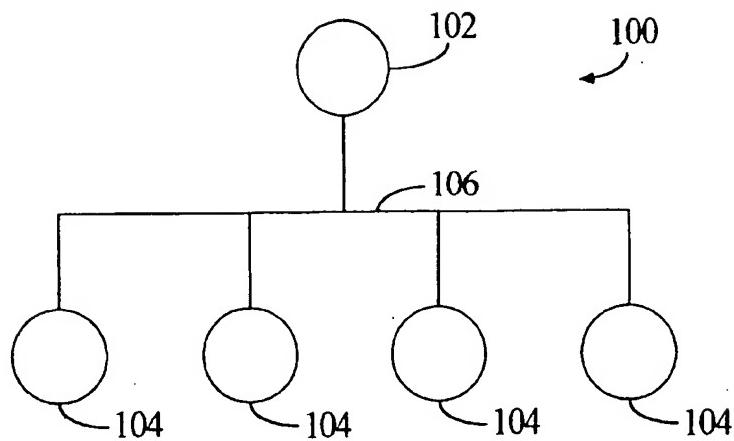


FIGURE 1

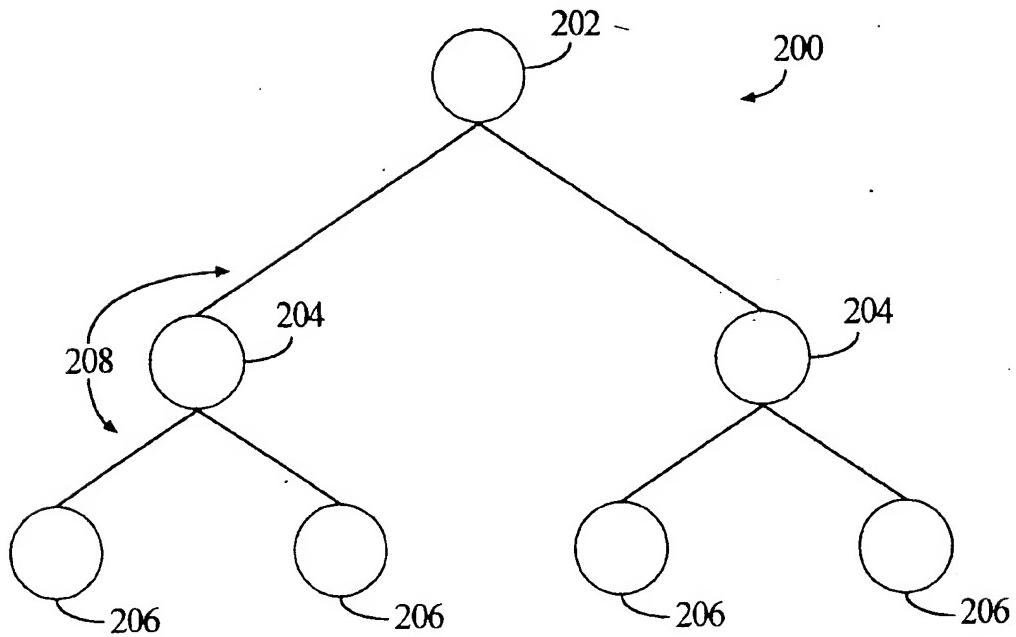


FIGURE 2

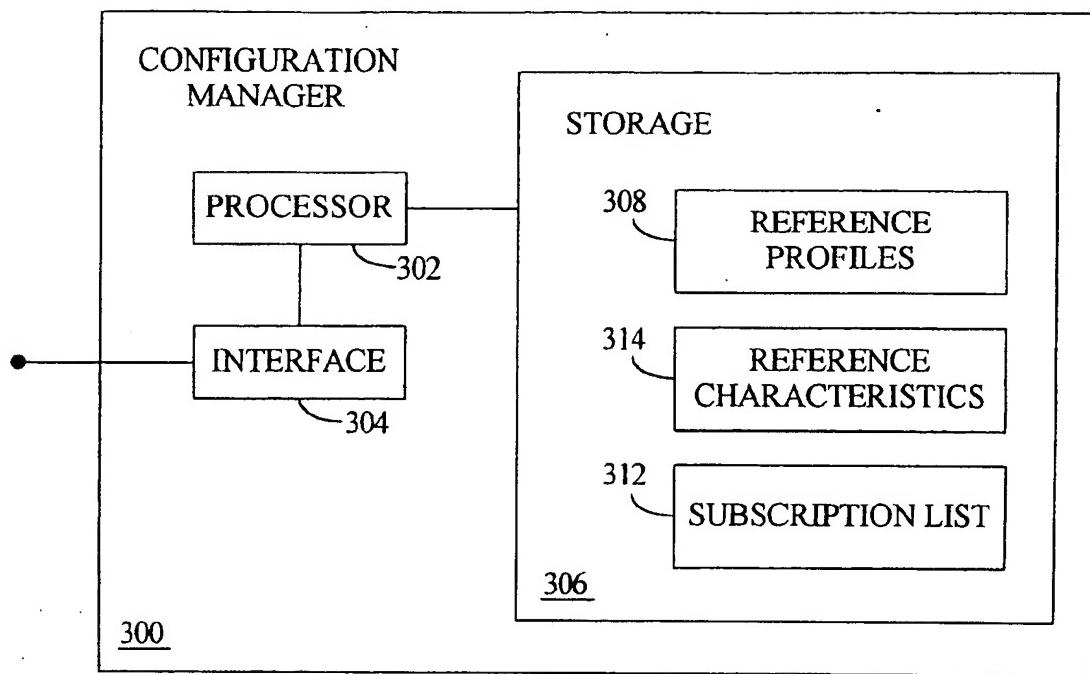


FIGURE 3

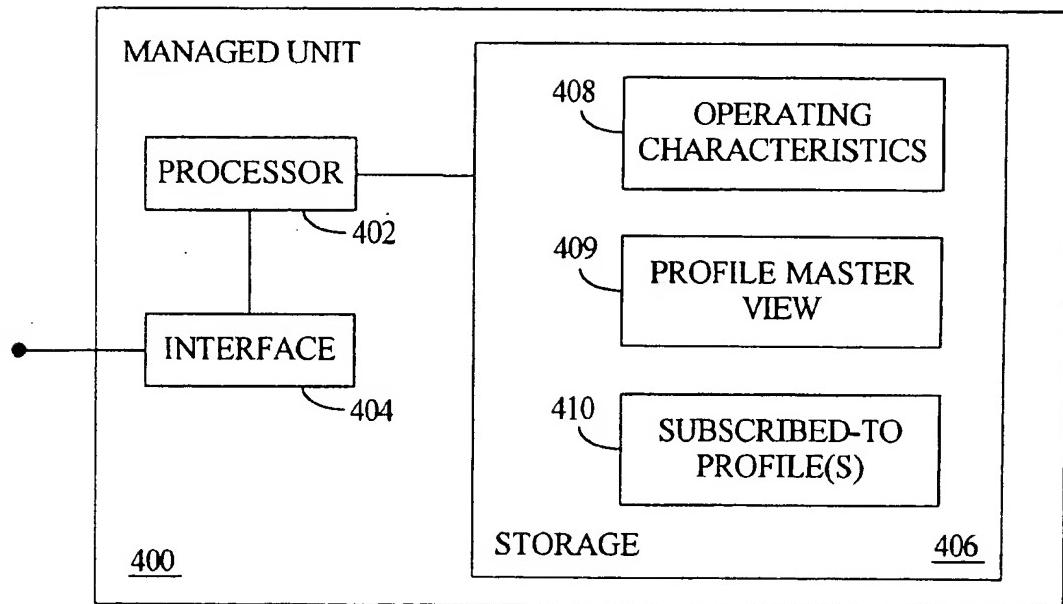


FIGURE 4

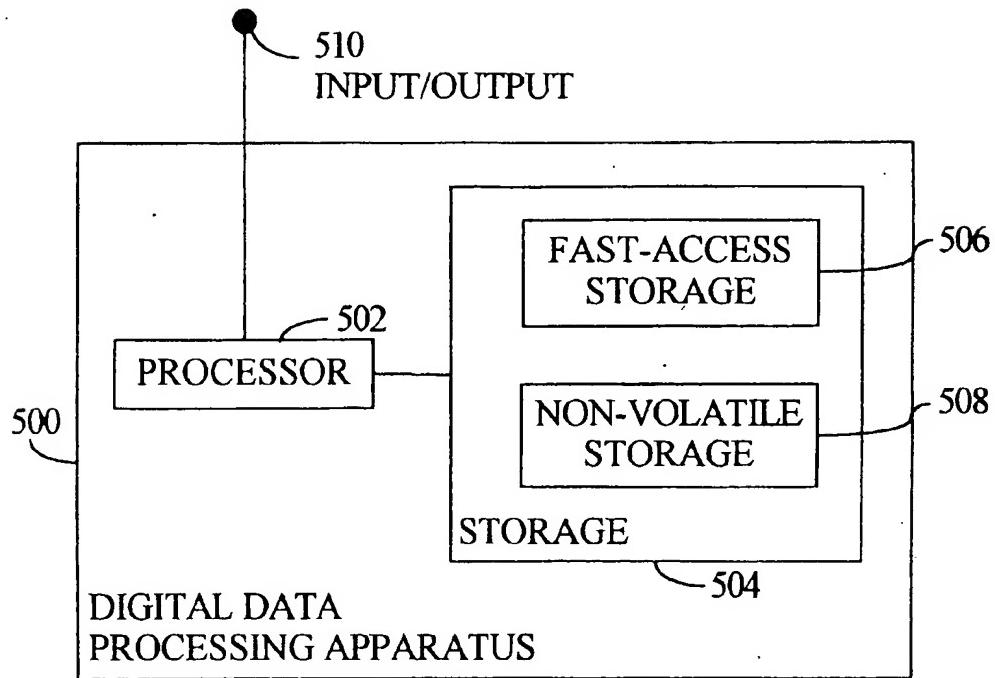


FIGURE 5

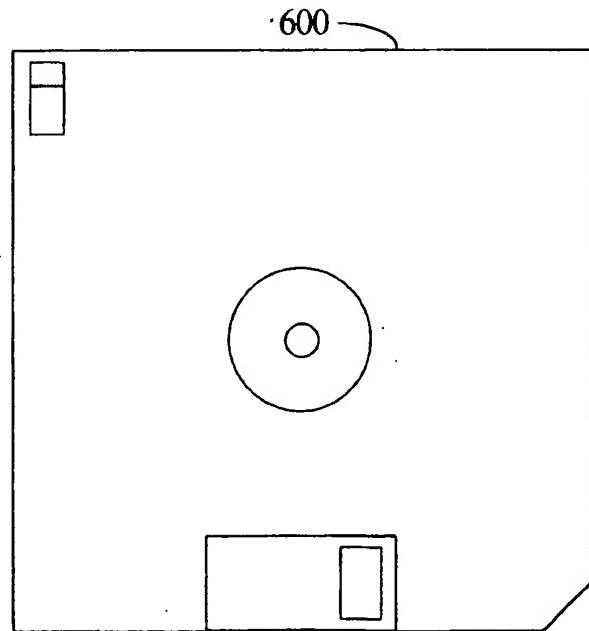


FIGURE 6

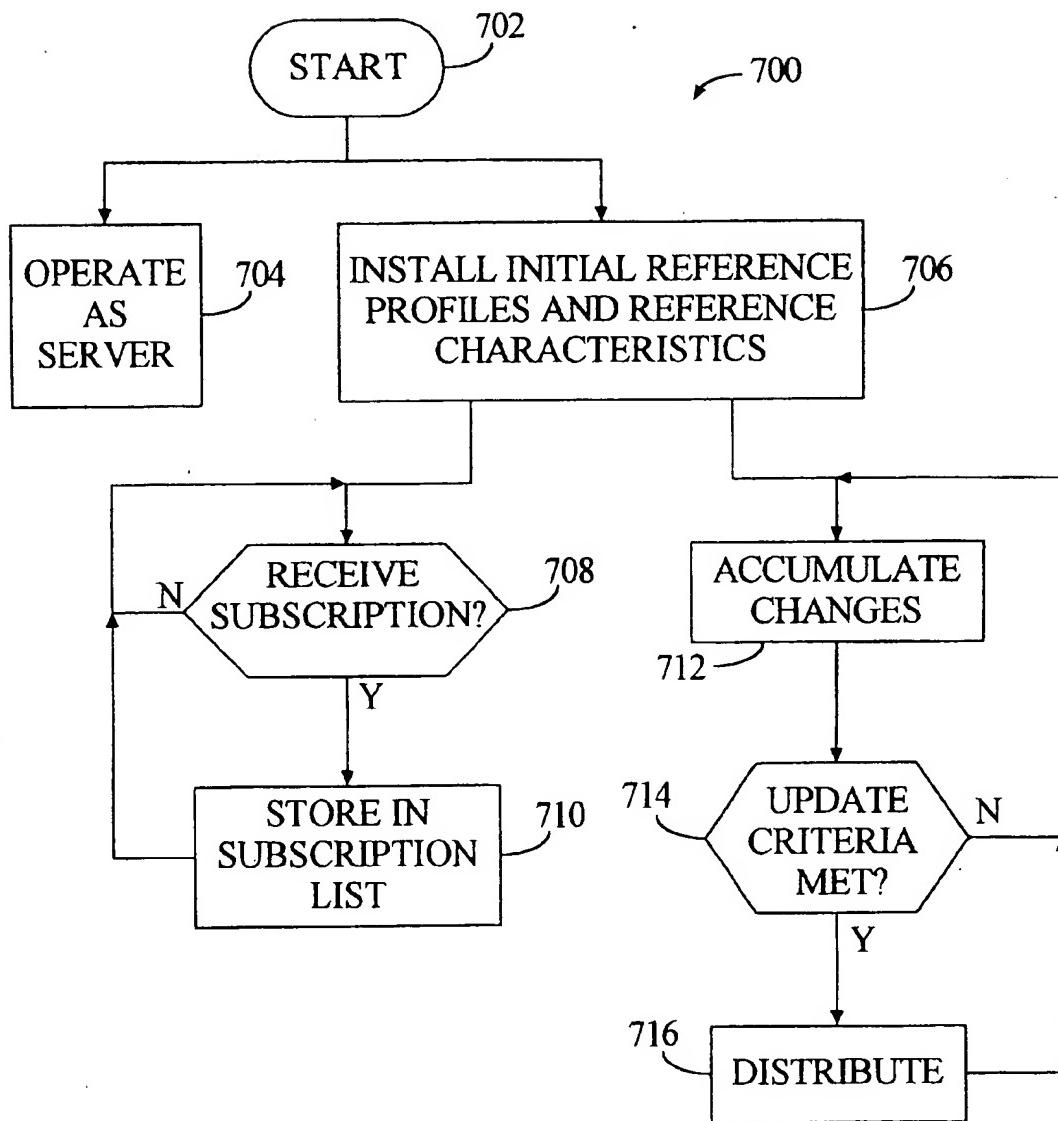


FIGURE 7

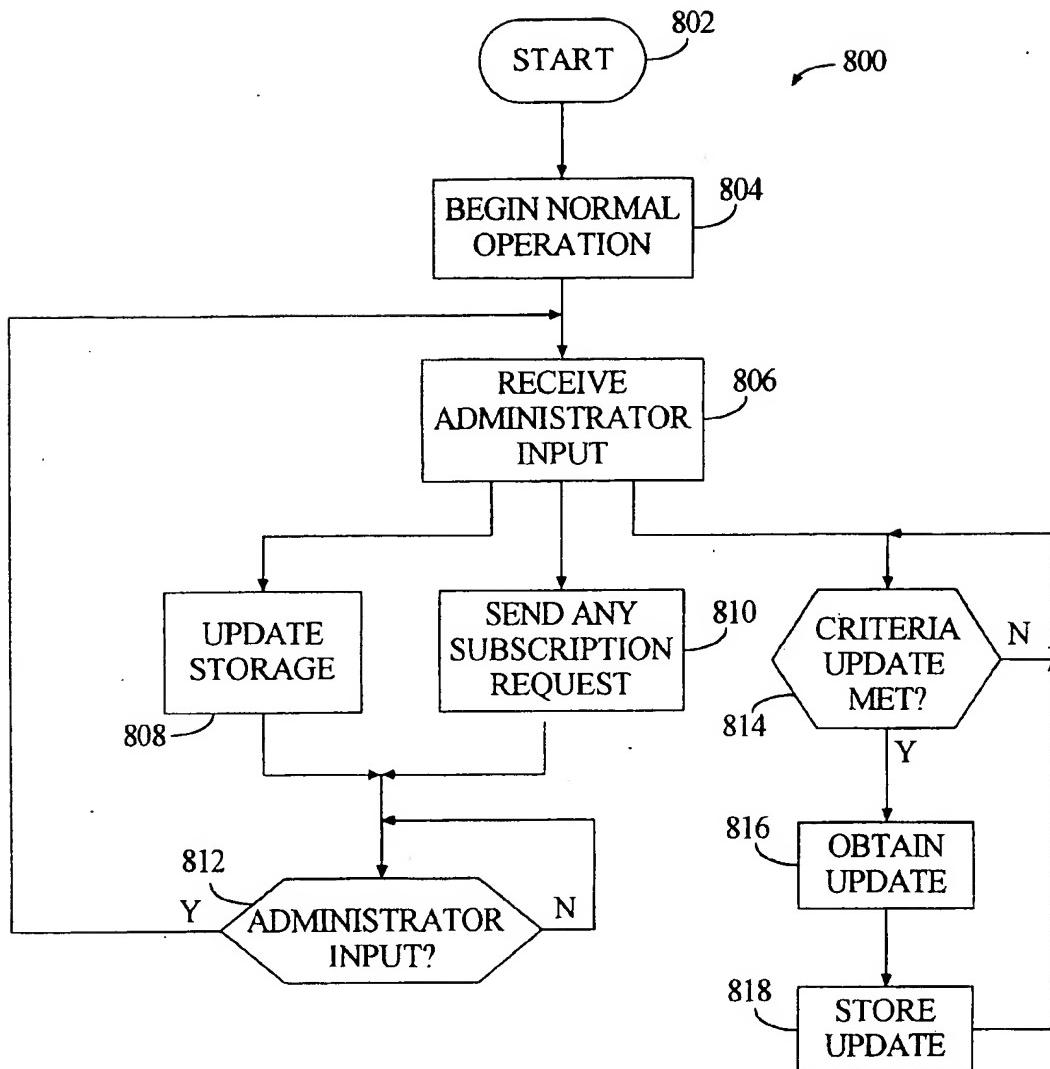


FIGURE 8

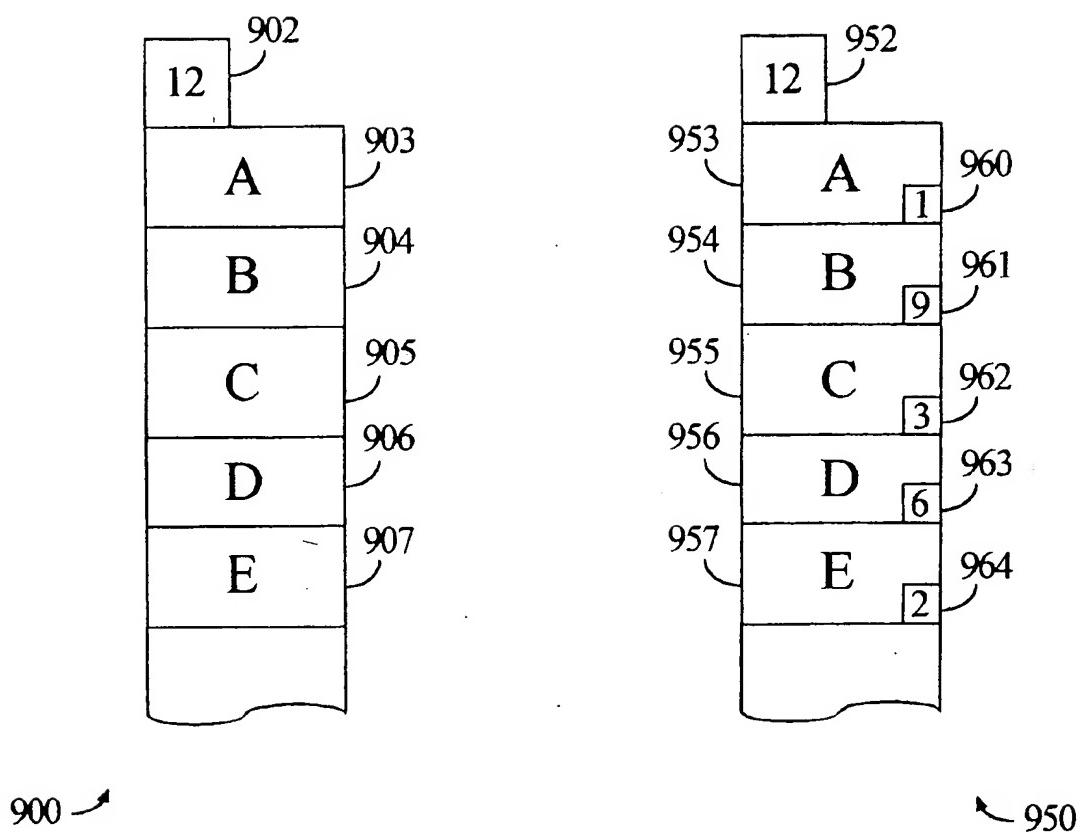


FIGURE 9

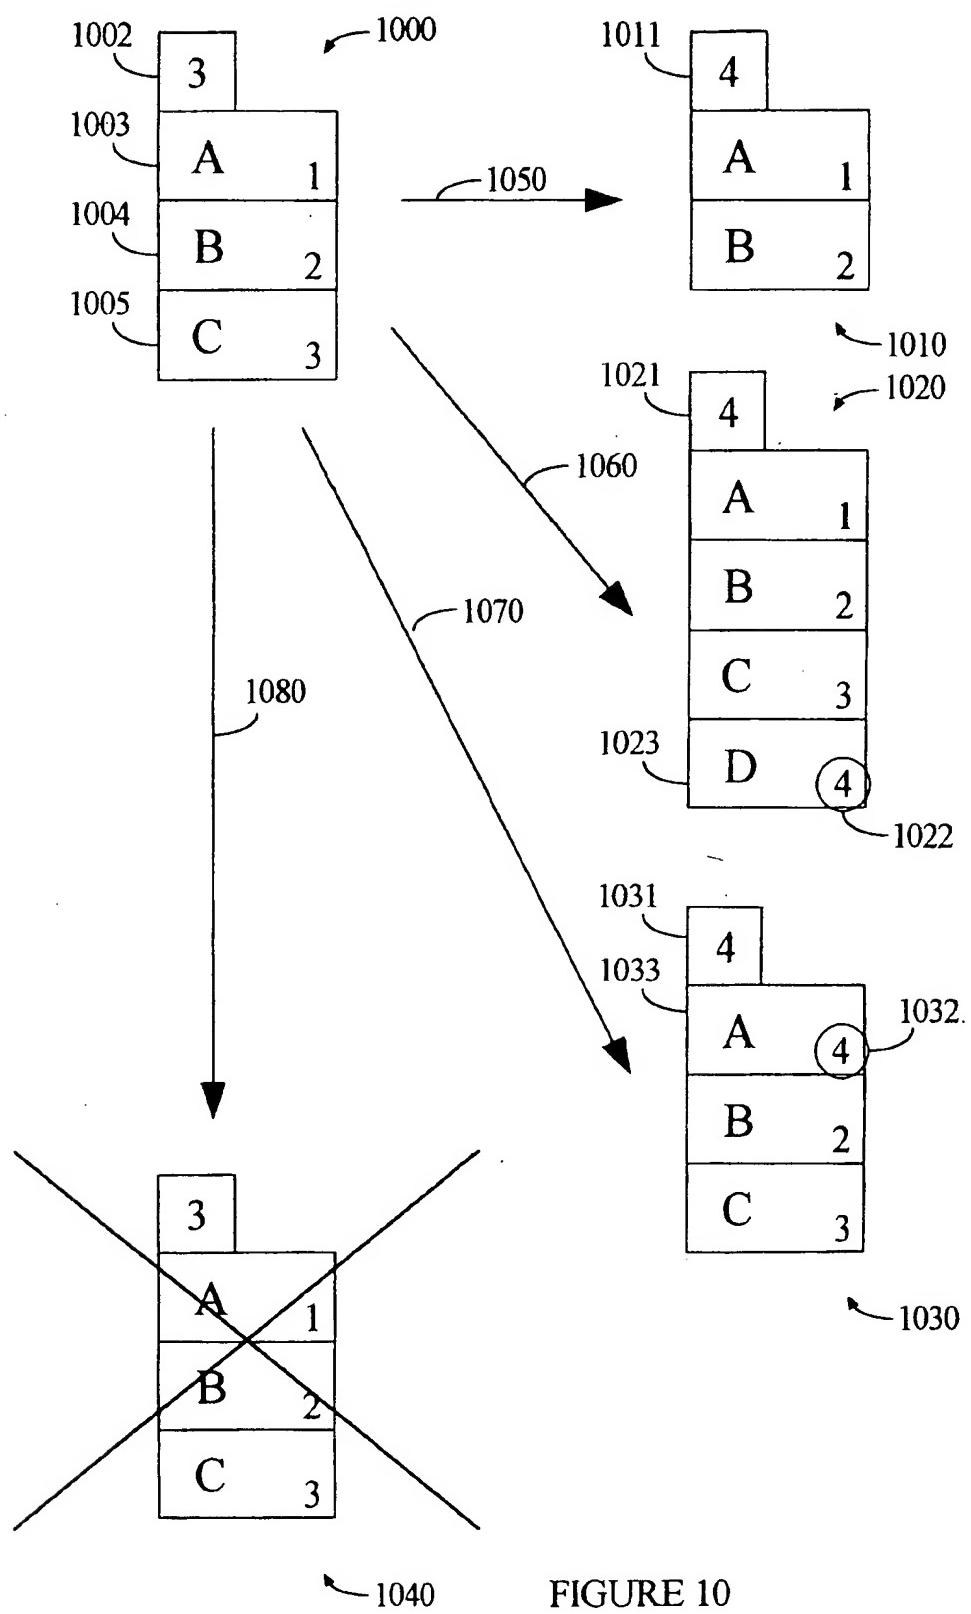


FIGURE 10

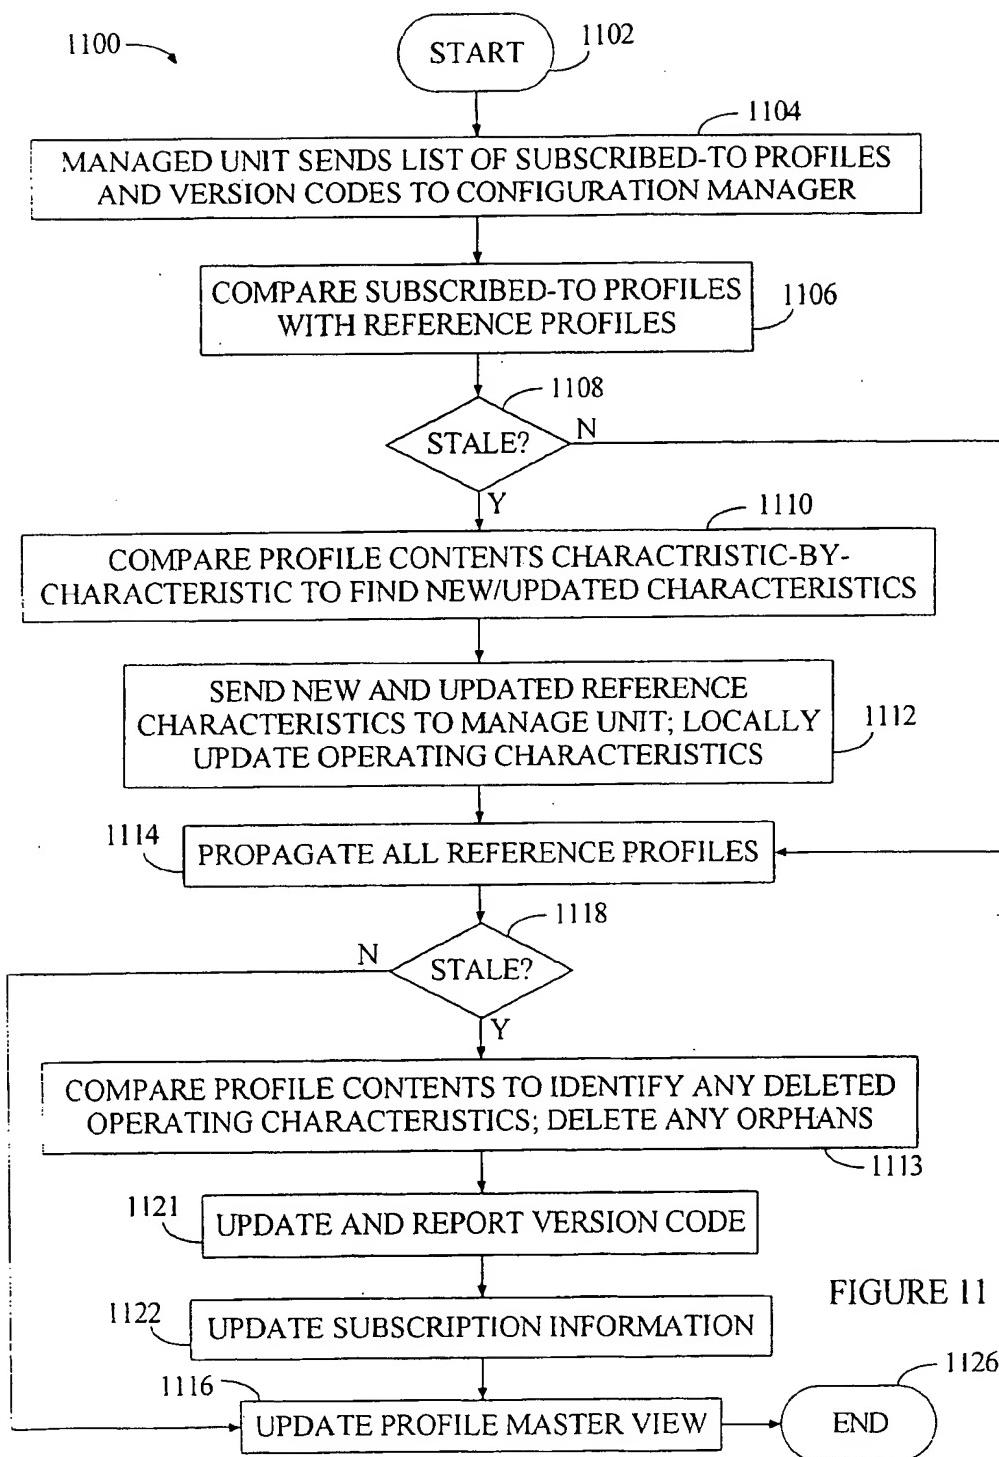


FIGURE 11

**COMPARATIVE UPDATES TRACKING TO  
SYNCHRONIZE LOCAL OPERATING  
PARAMETERS WITH CENTRALLY  
MAINTAINED REFERENCE PARAMETERS  
IN A MULTIPROCESSING SYSTEM**

**CROSS-REFERENCE TO RELATED  
APPLICATIONS**

The present application is related to co-pending U.S. application Ser. No. 09/144,424, entitled "Multiprocessing System with Automated Propagation of Changes to Centrally Maintained Configuration Settings," filed herewith, in the names of David Maxwell Cannon et al., and assigned to International Business Machines Corp. (IBM).

**BACKGROUND OF THE INVENTION**

**1. Field of the Invention**

The present invention relates to multiprocessing systems. More particularly, the invention concerns a multiprocessing system in which hierarchically superior configuration managers maintain various reference parameters that are selectively copied by subordinate managed units to form local operating parameters that subsequently guide the operation of these managed units. The invention includes a comparative technique for tracking updates to the reference parameters, and synchronizing each local operating parameter counterpart accordingly.

**2. Description of the Related Art**

In different forms, multiprocessing systems are in widespread use today. As one example, multiprocessing systems include "supercomputers" and other similar machines that utilize many separate high-speed processors. This arrangement is popular because, collectively, the processors provide a tremendously powerful computing engine. Multiprocessing systems also include "distributed processing systems," which use typically multiple physically distinct computing machines to cooperatively perform a task. With this arrangement a computing task that is too intensive for any one machine can be spread out and processed in smaller parts by many different machines. In addition to these examples, multiprocessing systems can include a hybrid of the two, or a variety of other arrangements with multiple computers, microprocessors, application program instances, program threads, or other processing elements.

In one popular example of distributed processing system, many servers are networked to provide a data storage system that provides users with many sophisticated data management capabilities. This arrangement may be called a distributed storage system; examples of this approach may be found in various known networks implementing the ADSTAR™ Distributed Storage Manager (ADSM)™ software of IBM. Although this and other similar systems constitute a significant advance and even enjoy widespread commercial success today, IBM continually seeks to improve the performance and efficiency of their multiprocessing systems. One area of focus concerns possible improvements to the ease and efficiency of configuring members of such multiprocessing systems.

Particularly, one drawback of the known arrangements is the time needed to individually configure each member of the multiprocessing system to operate as desired. In a system of networked servers, for example, configuration may involve a system administrator executing various configuration commands at a number of different networked servers to initially configure the servers, and thereby establish the

servers' operating characteristics. In many cases, this is done by the system administrator individually logging-in to each desired server and then executing a command. In more advanced systems, the system administrator may remotely log-in to desired servers remotely via another, nearby server. In either case, when many different servers are involved, there is a significant time penalty, since the administrator must take steps to execute each configuration command at the desired server. Moreover, this time penalty can arise repeatedly, since the same steps are required to reconfigure a server due to any updated operating requirements.

This time inefficiency presents a number of problems. First, it burdens the system administrator with work, making the system more costly to operate. As another limitation, the manual configuration and reconfiguration technique can be too error prone for some applications. In some cases, for example, many storage servers may be configured identically and distributed throughout a large office, for reasons of performance and scalability. In this situation, it is desirable to maintain the configuration of these servers in lock step, and thus all configuration changes must be universally applied to all servers. However, this is difficult and vulnerable to error due to the lengthy and repetitive entry of configuration commands at each server. Consequently, the known techniques for configuring members of a distributed processing system are not completely adequate for some applications due to certain unsolved problems.

**SUMMARY OF THE INVENTION**

Broadly, the present invention concerns a multiprocessing system in which hierarchically superior configuration managers maintain various reference parameters that are selectively copied by subordinate managed units to form local operating parameters that subsequently guide the operation of these managed units. The invention includes a comparative technique for tracking updates to the reference parameters, and synchronizing each local operating parameter counterpart accordingly.

The configuration manager maintains a list of reference profiles, and a list of reference characteristics. Each reference profile specifies one or more of the reference characteristics. Also in the configuration manager's storage, a subscription list names each managed unit subscribing to a reference profile, and the identity of that reference profile.

Each managed unit has a list of subscribed-to profiles and a list of operating characteristics, with each subscribed-to profile specifying one or more of the operating characteristics. To reduce network traffic and expedite update tracking, each managed unit may also maintain a local copy of the reference profiles, called a "master profile view." The operation of each managed unit is governed by its locally maintained operating characteristics.

When certain update criteria are satisfied, a managed unit and the configuration manager cooperatively synchronize the managed unit's operating characteristics with the configuration manager's reference characteristics. At this time, the managed unit's profile master view may also be synchronized with the configuration manager's list of reference profiles. The update criteria may trigger synchronization on any desired schedule of the managed unit and/or configuration manager, such as a periodic schedule, a non-periodic schedule, after a minimum number of updates occur, etc. Different managed units may have different update criteria, if desired.

To begin synchronization, the managed unit sends its list of subscribed-to profiles to the configuration manager. In an

alternative embodiment, the configuration manager sends its list of reference profiles to the managed unit. Then, these two lists are compared by the configuration manager (or by the managed unit in the alternative embodiment). This comparison is performed to identify any subscribed-to profiles whose counterparts in the list of reference profiles have been updated, and may be performed, for example, by comparing version codes. Such subscribed-to profiles are referred to as "stale." It is important to identify stale subscribed-to profiles because they might not properly identify current operating characteristics that the managed unit should be operating under.

For each stale subscribed-to profile, the stale subscribed-to profile is compared to its updated counterpart reference profile. This comparison identifies any "new" reference characteristics that are specified by the reference profile without having a counterpart operating characteristic specified by the stale subscribed-to profile. This comparison also identifies any updated reference characteristics, where an older version is named by the subscribed-to profile. In addition, this comparison identifies any reference characteristics that have been deleted from the reference profile but still remain in the corresponding subscribed-to profile. For each new or updated operating characteristic, the managed unit copies the counterpart reference characteristic from the configuration manager into the managed unit's list of operating characteristics. For each deleted operating characteristic, if the deleted reference characteristic is no longer named in any other subscribed-to profile, the operating characteristic is removed from the managed unit's list of operating characteristics.

Having synchronized the managed unit's list of operating characteristics, each stale subscribed-to profile is updated to match its up-to-date reference profile counterpart. This may be done, for example, by giving the subscribed-to profile the reference profile's version code. To complete synchronization, the managed unit's profile master view is updated to match the configuration manager's list of reference profiles.

Accordingly, one embodiment of the invention concerns a method for comparatively tracking updates to reference parameters, and synchronizing each local operating parameter counterpart accordingly. In another embodiment, the invention may be implemented to provide an apparatus such as a multiprocessing system, configuration manager, and/or managed unit, configured to operate as discussed above. In still another embodiment, the invention may be implemented to provide a signal-bearing medium tangibly embodying a program of machine-readable instructions executable by a digital data processing apparatus to perform a method to operate one or more components of a multiprocessing system, as discussed above.

The invention affords its users with a number of distinct advantages. First, the invention reduces the time needed to configure and reconfigure computers in a multiprocessing system. Instead of manually issuing commands to configure managed units, the managed units subscribe to centrally maintained configuration profiles and reference characteristics, and automatically receive any changes to the subscribed profiles and reference characteristics. Thus, instead of reconfiguring each managed unit, a system administrator enters changes once at the configuration manager; thereafter, these changes are automatically propagated down to all subscribing managed units. In addition to saving time, this reduces the potential for committing errors in the otherwise tedious work of repeatedly issuing the same configuration commands to the different nodes.

Additionally, the invention still maintains the flexibility for managed units to receive and implement local operating characteristics, apart from any subscriptions.

As a further benefit, the invention minimizes network traffic using version codes to identify new and updated reference characteristics, and then exchanging information as needed rather than all reference characteristics in a subscribed-to profile. Network traffic is also reduced by use of the profile master view, which facilitates profile comparison by the managed unit without having to communicate with the configuration manager. In addition, the invention avoids the need for centrally monitoring the application of updates at all affected managed units, and for managing the necessary history log. Instead, detailed comparison of a managed unit's subscribed-to profiles to the counterpart reference profiles quickly shows which changes (if any) the managed unit requires to its operating characteristics and/or profile master view. The invention also provides a number of other advantages and benefits, which should be apparent from the following description of the invention.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of one exemplary architecture of multiproCESSing system in accordance with the invention.

FIG. 2 is a block diagram of a different exemplary architecture of multiproCESSing system, configuration managers at multiple levels, in accordance with the invention.

FIG. 3 is a block diagram of an exemplary configuration manager in accordance with the invention.

FIG. 4 is a block diagram of an exemplary managed unit in accordance with the invention.

FIG. 5 is a block diagram of a digital data processing machine in accordance with the invention.

FIG. 6 shows an exemplary signal-bearing medium in accordance with the invention.

FIG. 7 is a flowchart of an exemplary operating sequence for a configuration manager in accordance with the invention.

FIG. 8 is a flowchart of an exemplary operating sequence for a managed unit in accordance with the invention.

FIG. 9 is a block diagram illustrating the components of an exemplary master profile and an exemplary reference profile, in accordance with the invention.

FIG. 10 is a block diagram showing how the components of a reference profile change with deletion, addition, and updating of a reference characteristic, as well as deletion of the reference profile entirely, in accordance with the invention.

FIG. 11 is a flowchart of an exemplary operating sequence for analyzing locally stored operating parameters against centrally maintained reference parameters, and locally implementing the resultant changes.

#### DETAILED DESCRIPTION

The nature, objectives, and advantages of the invention will become more apparent to those skilled in the art after considering the following detailed description in connection with the accompanying drawings. As mentioned above, the invention concerns a multiproCESSing system in which hierarchically superior configuration managers maintain various reference parameters that are selectively copied by subordinate managed units to form local operating parameters that subsequently guide the operation of these managed units. The invention includes a comparative technique for tracking

updates to the reference parameters, and synchronizing each total operating parameter counterpart accordingly.

#### HARDWARE COMPONENTS & INTERCONNECTIONS

##### System Architecture

One aspect of the invention concerns a multiprocessing system, which may be embodied by various hardware components and interconnections. The system includes multiple nodes, including one or more configuration managers and one or more managed units. Each managed unit is coupled to a hierarchically superior configuration manager. Configuration managers themselves may also act as managed units with respect to other, hierarchically even more superior configuration managers.

One example is the system 100 shown in FIG. 1. The system 100 includes multiple nodes, including a configuration manager 102 and multiple managed units 104. The managed units 104 and configuration manager 102 are connected by a communications link 106.

FIG. 2 depicts another example, with multiple managed units 206 that operate under a pair of configuration managers 204. The configuration managers 204, which answer to a higher-level configuration manager 202, may themselves operate as managed units in addition to their roles as configuration managers. The nodes of FIG. 2 are interconnected as shown by the communications link 208.

In addition to the arrangements of FIGS. 1-2, there are many other possible architectures embodying the invention, as will be apparent to those of ordinary skill in the art having the benefit of this disclosure. The multiprocessing system of the invention may be implemented with similar subcomponents, whether embodied by the system 100 (FIG. 1), the system 200 (FIG. 2), or another arrangement. As an example, each node (e.g., configuration manager or managed unit) of the system may comprise a processing element, such as a computing machine, software, or a combination. As examples of computing machines, each node may comprise a personal computer, mainframe computer, workstation, microprocessor, or another digital data processing machine. In contrast, a node may instead comprise an instance of an application program, a software thread, subroutine, or another software construct. Nodes may also be provided by combinations of hardware, software, firmware, and the like.

Each node, whether configuration manager or managed unit, may be implemented by similar or even identical hardware. Accordingly, the configuration managers and managed units may operate together on a "peer-to-peer" basis. As an example, all nodes may comprise digital data storage servers, such as IBM brand RS/6000 machines, running the ADSTAR Distributed Storage Manager (ADSM) storage management software.

The node-interconnecting communications link, such as the links 106 or 208, comprises a suitable hardware and/or software mechanism to relay messages between the nodes. For example, the links 106, 208 may comprise telephone lines, cable television lines, an Ethernet, token ring, ATM network, local area network, wide area network, Internet, Intranet, etc. In implementations where two or more nodes comprise separate application program instances on the same machine, the communications link also includes a motherboard, backplane, LAN card, bus, or other necessary hardware and/or software interconnecting these nodes. Communications over these links may employ any suitable protocol, such as TCP/IP, APPC, etc.

#### Configuration Manager

FIG. 3 shows the subcomponents of an exemplary configuration manager 300 in more detail. The configuration manager 300 includes a processor 302, an interface 304, and a storage 306. The processor 302 may comprise one or more microprocessors, an application-specific integrated circuit, or any other suitable digital data processing element. The interface 304 facilitates communications between the processor 302 and an attached communications link (not shown). The interface 304 may comprise a modem, intelligent input/output channel, bus, or any other suitable structure.

The storage 306 comprises writable non-volatile storage such as magnetic disk storage media, magnetic tape, optical media, battery-supported random access memory (RAM), and the like. The storage 306 includes reference profiles 308, reference characteristics 314, and a subscription list 312. The reference profiles 308 and reference characteristics 314 may be collectively called "reference parameters."

Generally, the reference characteristics 314 comprises a centralized storage of various machine-readable configuration settings. To all managed units of the configuration manager 300 ("subordinate" managed units), the reference characteristics 314 are read-only. As explained below, the subordinate managed units copy selected reference characteristics to local storage, and then use the copied characteristics as locally stored "operating characteristics." Each managed units uses its locally stored operating characteristics to guide its operation.

The list of reference profiles 308 includes a number of individual reference profiles, where each profile is a group of one or more reference characteristics. To the subordinate managed units, the reference profiles 308 are read-only. As discussed below, the managed units "subscribe" to selected reference profiles, thereby indicating a desire to keep locally stored operating characteristics up-to-date with the state of the selected reference profiles. The configuration manager 300 uses the subscription list 312 to track which subordinate managed units have subscribed to which reference profiles. The configuration manager 300 also uses the subscription list 312 to track whether the subscribing managed units have received up-to-date reference characteristics. This is facilitated by using version codes to distinguish updated profiles from earlier versions, as discussed in greater detail below.

To provide a more concrete example, Tables 1-3 (below) depict an exemplary set of reference profiles 308, reference characteristics 314, and subscription list 312 in a data storage system where each node is a data storage server.

TABLE 1

Reference Profiles	
PROFILE NAME	ASSOCIATED REFERENCE CHARACTERISTICS
Profile A	1, 3, 4, 8
Profile B	1, 2
Profile C	1
Profile D	5

TABLE 2

Reference Characteristics	
REFERENCE CHARACTERISTIC NUMBER	DETAILS
1	ADMINISTRATOR=MIKE (Password=Wildcat, Authority=System)
2	ADMINISTRATOR=DAVE (Password=Cougar, Authority=System)
3	ADMINISTRATOR=CAROLYN (Password=Penguin, Authority=Policy)
4	SCHEDULE=INCR_BACKUP (Frequency=Daily, Begintime=02:00, DayOfWeek=Any, Action=Incremental)
5	SCHEDULE=WEEKLY_ARCHIVE (Frequency=Weekly, Begintime=18:30, DayOfWeek=Saturday, Action=Archive)
6	SCHEDULE=DB_BACKUP (Frequency=Daily, Begintime=05:00, DayOfWeek=Any, Action=DbBackup)
7	SERVER=CONFIG_MGR (Address=9.115.12.92 Port=1500, Password=batman)
8	SERVER=TUCSON_BRANCH (Address=9.115.39.123, Port=1500, Password=robin)

20

TABLE 3

Subscription List	
MANAGED UNIT	SUBSCRIBED-TO PROFILES
DEVICE NUMBER 0001	Profile A (version code 1)
DEVICE NUMBER 0002	Profile A (version code 1)
DEVICE NUMBER 0003	Profile A (version code 1), Profile D (version code 6)
DEVICE NUMBER 0004	Profile B (version code 9), Profile C (version code 3)
DEVICE NUMBER 0005	Profile D (version code 6)

## Managed Unit

FIG. 4 shows the subcomponents of an exemplary managed unit 400 in more detail. The managed unit 400 includes a processor 402, an interface 404, and a storage 406. The processor 402 may include one or more microprocessors, an application-specific integrated circuit, or any other suitable digital data processing element. The interface 404 facilitates communications between the processor 402 and an attached communications link (not shown). The interface 404 may comprise a modem, intelligent input/output channel, bus, or any other suitable structure.

The storage 406 comprises writable non-volatile storage such as magnetic disk storage media, magnetic tape, optical media, battery-supported random access memory (RAM), and the like. The storage 406 includes operating characteristics 408, subscribed-to profiles 410, and a profile master view 409.

Generally, the subscribed-to profiles 410 contain the names of zero, one, or more reference profiles 308 (FIG. 3), to which the managed unit 400 subscribes. Each subscribed-to profile also includes a version code, as explained below. The content of the reference profiles to which the managed unit 400 subscribes are contained in the profile master view 409, along with the content of all reference profiles 308, as discussed below. The significance of subscription for the managed unit 400 is that the managed unit 400 maintains its own copy of all reference characteristics 314 specified by its subscribed-to profiles 410. This local copy of the reference characteristics 314 is embodied by the operating characteristics 408. The managed unit 400 operates in compliance with the locally stored operating characteristics 408. The operating characteristics 408 may also include other opera-

tional characteristics unrelated to the subscribed-to profiles. These additional characteristics, called "local" operating characteristics, may originate from a system administrator, local application program, or another source.

25 Since a configuration manager's centrally stored reference characteristics 314 and reference profiles 308 may be changed from time to time, the invention provides means to distribute the changes among all affected subordinate managed units. This technique is discussed in greater detail below.

30 The profile master view 409 is a copy of the configuration manager's reference profiles 308, which the managed unit employs in various ways as discussed below. As explained below, the profile master view may be used to query profile information without any network overhead, i.e., communication with the configuration manager. Another benefit of locally storing the profile master view is that managed units can identify which profiles or associated characteristics have been deleted since the last refresh by comparing profile contents. The configuration manager does not need to maintain historical information such as lists of deleted operational characteristics.

## Exemplary Digital Data Processing Apparatus

35 Another aspect of the invention concerns a digital data processing apparatus, provided to implement one or more managed units or configuration manager. This apparatus may be embodied by various hardware components and interconnections, as discussed below.

40 FIG. 5 shows an example of one digital data processing apparatus 500. The apparatus 500 includes a processor 502, such as a microprocessor or other processing machine, coupled to a storage 504. In the present example, the storage 504 includes a fast-access storage 506, as well as nonvolatile storage 508. The fast-access storage 506 may comprise RAM, and may be used to store the programming instructions executed by the processor 502. The nonvolatile storage 508 may comprise, for example, one or more magnetic data storage disks such as a "hard drive," a tape drive, or any other suitable storage device. The apparatus 500 also includes an input/output 510, such as a line, bus, cable, electromagnetic link, or other means for the processor 502 to exchange data with other hardware external to the apparatus.

45 Despite the specific foregoing description, ordinarily skilled artisans (having the benefit of this disclosure) will

recognize that the apparatus discussed above may be implemented in a machine of different construction, without departing from the scope of the invention. As a specific example, one of the components 506, 508 may be eliminated; furthermore, the storage 504 may be provided on-board the processor 502, or even provided externally to the apparatus 500.

### OPERATION

In addition to the various hardware embodiments described above, a different aspect of the invention concerns a method for operating a multiprocessing system in which hierarchically superior configuration managers maintain various reference parameters that are selectively copied by subordinate managed units to form local operating parameters that subsequently guide the operation of these managed units. The invention includes a comparative technique for tracking updates to the reference parameters, and synchronizing each local operating parameter counterpart accordingly.

#### Signal-Bearing Media

In the context of FIGS. 1-5, such a method may be implemented, for example, by operating the configuration managers and managed units, as embodied by digital data processing apparatuses 500, to execute respective sequences of machine-readable instructions. These instructions may reside in various types of signal-bearing media. In this respect, one aspect of the present invention concerns a programmed product, comprising signal-bearing media tangibly embodying a program of machine-readable instructions executable by a digital data processor to perform the functions of configuration manager or managed unit as explained herein.

This signal-bearing media may comprise, for example, RAM (not shown) contained within storage of the configuration manager or managed unit, as represented by the fast-access storage 506 for example. Alternatively, the instructions may be contained in another signal-bearing media, such as a magnetic data storage diskette 600 (FIG. 6), directly or indirectly accessible by the processor 502. Whether contained in the storage 506, diskette 600, or elsewhere, the instructions may be stored on a variety of machine-readable data storage media, such as direct access storage (e.g., a conventional "hard drive" or a RAID array), magnetic tape, electronic read-only memory (e.g., ROM, EPROM, or EEPROM), an optical storage device (e.g. CD-ROM, WORM, DVD, digital optical tape), paper "punch" cards, or other suitable signal-bearing media including transmission media such as digital and analog and communication links and wireless. In an illustrative embodiment of the invention, the machine-readable instructions may comprise software object code, compiled from a language such as "C," etc.

#### Configuration Manager: Operation

As mentioned above, configuration managers maintain various reference profiles to which subordinate managed units selectively subscribe. Whenever the profiles or their reference characteristics change, the configuration managers propagate the changes down to all affected managed units. FIG. 7 shows an exemplary process sequence 700 to describe the operation of an exemplary configuration manager in a multiprocessing system according to the invention. For ease of explanation, but without any intended limitation, the example of FIG. 7 is described in the context of the hardware environment described above in FIGS. 1-4.

The operations 700 are initiated in step 702, which may be performed when the configuration manager 300 is powered-up, re-started, or otherwise begins operation. Following step 702, the configuration manager 300 may (optionally) conduct other operations 704, aside from its role as a configuration manager. For instance, in the data storage system embodiment used presently, the configuration manager 300 may operate as a data server, as shown by step 704.

Concurrently with these non-manager operations (step 704), the configuration manager 300 performs tasks involved in overseeing the managed units (steps 706-716). Specifically, in step 706 the configuration manager installs an initial set of reference profiles 308 and reference characteristics 314. As an example, these initial settings may be manually set by a system administrator.

After step 706, the configuration manager 300 begins the concurrent processes of (1) receiving and processing managed units' subscriptions, and (2) accumulating and propagating changes to reference profiles and/or reference characteristics that affect subscribing managed units. Step 708 starts the subscription process. Namely, in step 708 the configuration manager 300 waits until it receives a subscription message from a managed unit. Subscription messages comprise requests from managed units to subscribe to one of the pre-defined reference profiles 308. As such, a subscription includes the name of the desired reference profile, and the identity of the managed unit requesting subscription. When the configuration manager 300 receives a subscription message in step 708, the configuration manager 300 updates the subscription list 312 to show the subscribed-to profile in association with the subscribing managed unit (step 710). Following step 710, step 708 waits for another subscription.

In parallel with steps 708-710, the configuration manager 300 collects and distributes changes to the reference parameters. As an example, these changes may originate from a system administrator (not shown) accessing the configuration manager. To provide another example, these changes may originate from another configuration manager that is hierarchically superior to the present configuration manager 300, where this superior configuration manager propagates the changes downward to the present configuration manager 300.

The configuration manager 300 receives and accumulates the changes in step 712. These changes may include changes to the reference characteristics 314 and/or the reference profiles 308. The accumulation of the changes in step 712 may involve (1) buffering changes and then implementing them in storage 306, (2) not buffering any changes but marking the changed profiles and/or characteristics in a directory such as a table, list, bitmap, etc., or (3) another approach.

In step 714, the configuration manager 300 determines whether certain prescribed update criteria have been met to start distributing the accumulated changes. The update criteria may be set by a system administrator or other user, permanently incorporated into the firmware or other programming of the processor 302 upon manufacture of the configuration manager 300, or established in another way. As mentioned above, the update criteria determine when the configuration manager 300 propagates changes in the reference parameters down to the subordinate managed servers. The criteria may take various embodiments in accordance with this invention, such as:

**Threshold Number of Changes:** As one example, the configuration manager 300 may deem the criteria met whenever it receives a threshold number of changes, such as five changes.

**Fixed Time Period:** As another example, the configuration manager 300 may deem the criteria met on a periodic basis, such as one hour. In this embodiment, the criteria may be met, for example, each time a repeating timer expires.

**Request by Subordinate Managed Units:** As still another example, the configuration manager 300 may deem the criteria met whenever one or more subordinate managed units submit requests to download accumulated changes. The managed units may submit such requests on any appropriate basis, such as periodically.

If step 714 finds that the update criteria are not met, step 712 continues to accumulate incoming changes. When the update criteria of step 714 are finally met, the configuration manager 300 in step 716 distributes or "propagates" the accumulated changes. To distribute the changes, the configuration manager 300 cross-references the accumulated changes with the subscription list 312 to identify all managed units requiring update propagation. If an accumulated change concerns a change to a reference profile, the configuration manager transmits the changed profile to all managed units, as discussed in greater detail below.

According to one embodiment of the invention, update criteria may be chosen so that they are met individually on a managed-unit-by-managed-unit basis, or more universally for all managed units together. Therefore, in step 716 the configuration manager 300 may propagate updates to all subscribing managed units, or to certain individual managed units requiring updates, depending on which update criteria are used.

If a reference profile has changed the configuration manager 300 also examines the changed profile to determine whether the profile includes any reference characteristics that were not previously in that profile. If there have been any added reference characteristics the configuration manager 300 also transmits these new reference characteristics to the subscribing managed units. As an additional feature, the configuration manager 300 may determine whether subscribing managed units already have such reference characteristics due to a subscription to a different profile; in this event, the configuration manager 300 may skip sending the duplicate reference characteristics. When the accumulated changes have been distributed, completing step 716, the routine 700 returns to step 712 to accumulate more changes.

#### Managed Unit: Operation

As mentioned above, managed units selectively subscribe to various reference profiles of constituent reference characteristics maintained by hierarchically superior configuration managers. Whenever the reference profiles or their reference characteristics change, the configuration managers automatically propagate the changes down to all managed units, which store and implement and then operate pursuant to the changes. FIG. 8 shows an exemplary process sequence 800 to describe the operation of an exemplary managed unit in a multiprocessing system according to the invention. For ease of explanation, but without any intended limitation, the example of FIG. 8 is described in the context of the hardware environment described above in FIGS. 1-4.

The operations are initiated in step 802, which may be performed when the managed unit 400 is powered-up, re-started, or otherwise begins operation. Following step 802, the managed unit 400 starts normal operations (step 804). In the illustrated example, where the managed unit is a storage server, these operations may involve satisfying storage access requests of one or more hosts, clients, or other

machines (not shown) coupled to the managed unit 400. While these operations continue, the managed unit 400 receives certain input in step 806, for the purpose of initially configuring the managed unit 400. In the illustrated example, this input is received from a system administrator (not shown), who accesses the configuration manager 400 via an attached console, remote terminal, another computer, or another suitable mechanism (not shown).

This initial administrator input may include manual entry of one or more locally-set operating characteristics, such as those configuration settings needed to enable communication with the configuration manager, perform basic local operations, etc. During subsequent post-initial performance of step 806, the administrator input may include items such as the following:

**New Subscription:** The input may include a request to subscribe to a profile, to which the managed unit 400 does not already subscribe.

**Deleted Subscription:** The input may specify deletion of a subscription to a profile that the managed unit subscribes to.

**Local Operating Characteristics:** The input may also include addition or deletion of one or more "locally-set" operating characteristics, added by manual input, rather than the managed unit receiving them via subscription. As a result, these operating characteristics are not updated by subscription. As on optional embodiment, the settings of local operating characteristics may automatically yield to any conflicting operating characteristics arising from subscribed-to profiles. This conflict resolution may be set by default, by administration election, or another appropriate means.

After the initial input is received in step 806, several things occur concurrently. Namely, the managed unit 400 updates the storage 406 according to the new input (step 808). In the case of request to add or delete subscribed-to profiles, the managed unit 400 adds (or deletes) the profiles to (or from) the list of subscribed-to profiles 410. In the case of a locally-set operating characteristics, the managed unit 400 adds (or deletes) the operating characteristics from the list 408. After step 808, the configuration manager 400 determines whether it has received any more input (step 812), such as administrator input in the present example. When the managed unit 400 receives further input, step 812 returns to step 806 to receive this input, as discussed above.

Concurrently with step 808, the managed unit 400 sends any necessary subscription requests (step 810), in order to implement the input of step 806. For instance, if the step 806 input requested subscription to a new profile, the managed unit 400 in step 810 sends a subscription message to the configuration manager 300 to subscribe to that profile. The subscription message includes an identification of the profile and the managed unit 400 requesting subscription. If the input of step 806 requested removal of a subscribed-to profile, the managed unit 400 may send an appropriate message to the configuration manager 300 requesting withdrawal from subscription to that profile. Alternatively, the managed unit 400 may add or delete profiles to/from the list of subscribed-to profiles 410, and wait for the configuration manager 300 to initiate updating of its subscription list 312 to match the managed units' lists of subscribed-to profiles 410. In either case, processing of subscriptions by the configuration manager 400 is discussed above in step 708, FIG. 7.

Concurrently with steps 808, 810, and 812, the managed unit 400 determines whether the prescribed criteria have been met (step 814) to obtain reference parameter updates

accumulated by the configuration manager 300. The update criteria may be set by a system administrator or other user, permanently incorporated into the firmware or other programming of the processor 402 upon manufacture of the managed unit 400, or established in another way. As mentioned above, these criteria determine when the managed unit 400 obtains changes in the reference parameters from the configuration manager 300. The criteria may take various embodiments in accordance with this invention, such as:

**Locally-Determined Criteria:** In one example, the managed unit may deem the criteria met upon any appropriate local condition, such as (1) expiration of a periodic local timer, (2) according to a non-periodic local schedule, or (3) another basis.

**Propagation By Configuration Manager:** In another example, the configuration manager alone may determine when to propagate changes to subordinate managed units. As mentioned above, this decision may be made on the basis of receiving a threshold number of changes, expiration of a fixed time period, etc. In this embodiment, the managed unit 400 deems the update criteria to be met (step 814) when changes are propagated from the configuration manager.

Step 814 repeats until the update criteria are met, and then progresses to step 816. In step 816, the managed unit obtains the updated reference parameters. If the update criteria were met by local criteria (e.g., local timer expiration, prescribed local schedule, etc.), then step 816 involves the managed unit 400 requesting and receiving updated reference parameters (if any) from the configuration manager 300. In contrast, if the update criteria were met by the configuration manager propagating changes according to its own schedule, then step 816 involves the managed unit 400 receiving the propagated changes sent by the configuration manager 300. In either case, each update of steps 816-818 always includes a change to a master profile, and may also include changes to the operating characteristics 408.

In either of the foregoing embodiments, step 816 is followed by step 818, where the managed unit 400 stores the received updates in the storage 406. Following step 818, the routine 800 returns to step 814, awaiting satisfaction of the update criteria again.

#### Partial Reference Characteristic Propagation

As illustrated, when changes to a reference characteristic are propagated (step 716, FIG. 7), the entire content of that reference characteristic is transmitted. Alternatively, the configuration manager 300 may limit propagation to the changed parts of the reference characteristic. This saves transmission time, and expedites the update process. To support partial reference characteristic propagation, the configuration manager 300 is programmed to track the contents of the reference characteristics involved in each managed unit's subscriptions. Then, before distributing updates (step 716, FIG. 7), the configuration manager 300 compares the tracked contents to the current contents to specifically identify the changes. As one example, this process may be achieved by the configuration manager 300 logging changes to reference characteristics in a change log.

As a further expansion to reference characteristic propagation, it may be useful to permit hierarchically arranged reference characteristics. In this embodiment, a "parent" reference characteristic may include, among other things, one or more lower-level reference characteristics. When any lower-level reference characteristic is changed, the "parent" reference characteristic is also considered to be

changed. Changes to lower-level reference characteristics are accumulated in step 712, resulting in the subsequent distribution of the entire parent reference characteristic, in the manner discussed above. As an alternative, the propagation of changes to lower-level reference characteristics of a parent reference characteristic may be limited to the changed lower characteristics only, rather than the entire parent.

#### COMPARATIVE UPDATE TRACKING

##### Characteristic & Profile Versions

###### Introduction

FIG. 9 shows the components of an exemplary master profile 900 maintained by a managed unit (in the profile master view 409), and an exemplary reference profile 950 maintained by the configuration manager (in the reference profiles 308). Although other arrangements may be used without departing from the invention, the present example helps to illustrate one embodiment of the invention, and also provides a useful foundation for the comparative update tracking process, which is described below.

The reference profile 950 includes multiple reference characteristics (953-957) and a profile version code (952). The profile version code 952 is incremented once whenever the profile 950 is manually changed. Such manual changes may occur, for example, when the profile 950 is changed by a system administrator. The profile version code 952 is also incremented by one wherever any of the profile's constituent reference characteristics 953-957 experience an update.

The characteristic version codes 960-964 of the individual reference characteristics 953-957 represent the reference profile's version code at the respective times when the individual reference characteristics 953-957 were last added or updated. For example, the reference characteristic 953 was added upon the adoption of version number one of the profile 950, as shown by the characteristic version code 960. Similarly, the reference characteristic 956 was added upon the adoption of version number six of the profile 950, as shown by the characteristic version code 963.

###### Versions—Master Profile

Similar to the reference profile 950, each profile in the profile master view 409 includes multiple operating characteristics and a profile version code. An example is shown by the master profile 900, which includes multiple operating characteristics 903-907 and a profile version code 902. As explained below (step 1116, FIG. 11), each master profile is updated to match its counterpart reference profile, which serves as the source for that master profile. Specifically, the version code of a master profile is copied exactly from the version code of its reference profile counterpart. For example, if a master profile is updated according to reference profile version twelve, that master profile will assume version twelve itself.

###### Versions—Subscribed-to Profile

Although not shown in FIG. 9, each profile in the list of subscribed-to profiles 410 contains the name and version code of a reference profile that the managed unit 400 subscribes to. In the illustrated embodiment, subscribed-to profiles do not list their constituent operating characteristics. These are shown by the expanded profile contained in the profile master view 409, as shown above.

###### Identifiers

Rather than the actual content of the constituent operating characteristics themselves, the master profile 900 may list suitable identifiers 903-907 of the constituent reference

characteristics. Similarly, rather than the content of the constituent reference characteristics themselves, the reference profile 950 may list suitable identifiers 953-957 of the corresponding reference characteristics. In one embodiment, These identifiers comprise reference characteristic names; however, the identifiers may comprise table or other indices, address pointers, or other suitable identifiers instead.

The reference profiles 308 identify (e.g., by name) corresponding reference characteristics 314, all of which are stored together in the configuration manager's storage 306. In contrast, the constituent reference characteristics of the master profiles 409 may not all reside in the managed units storage 406; those reference characteristics named by subscribed-to profiles are stored in the list of operating characteristics 408, but the remaining referenced characteristics named by the master profiles 409 are stored only back at the configuration manager's list 314.

#### Changing a Reference Profile

FIG. 10 shows several illustrative update operations occurring to an exemplary reference profile 1000 stored in the list 308 at the configuration manager 300. Reference profiles may be changed under various circumstances, such as (1) a system administrator accessing the configuration manager, (2) downloading of changes from another, hierarchically superior, configuration manager, or (3) other situations.

In this example, the profile 1000 has a profile version code 1002 of three, and includes constituent reference characteristics 1003-1005. The transition 1050 shows how the reference profile 1000 in the list 308 is updated by deleting a constituent reference characteristic. Namely, the reference characteristic 1005 is deleted, yielding an updated reference profile 1011. Since the original reference profile 1000 had a profile version code 1002 of three, the new reference profile 1010 has a profile version code 1011 of four.

As an alternative to the transition 1050, the reference profile 1000 may instead be modified by adding another constituent reference characteristic as shown by the transition 1060. Namely, the reference characteristic 1023 is added, yielding an updated profile 1020. Since the original profile 1000 had a profile version code 1002 of three, the new reference profile 1020 has a profile version code 1021 of four. Moreover, the added reference characteristic 1023 shows a characteristics version code 1022 of four, since it was added upon adoption of the reference profile having a profile version code of four.

As an alternative to the transitions 1050, 1060, the reference profile 1000 may instead be modified by updating one of its constituent reference characteristics, as shown by the transition 1070. Namely, the reference characteristic 1003 is updated by introduction of a newer version of the same reference characteristic. In this example, the reference characteristic 1003 is replaced by the newer reference characteristic 1033. This yields an updated reference profile 1030. Since the original reference profile 1000 had a profile version code 1002 of three, the new reference profile 1030 has a profile version code 1031 of four. Moreover, the updated reference characteristic 1033 shows a characteristic version code 1032 of four, since it was added at the adoption of the reference profile having a profile version code of four.

As an alternative to the transitions 1050, 1060, 1070, the reference profile 11300 may instead be entirely deleted. In this case, the reference profile 1000 is removed from the list 308 by deleting its profile version number 1002, the constituent reference characteristic identifiers 1003-1005, and

the characteristic version codes. As explained below, deletion of the profile 1000 does not delete the constituent reference characteristics themselves, only the identifiers (e.g., pointers) present in the profile 1000.

#### Comparative Update Tracking: Operating Sequence Introduction

The foregoing illustrations show the structure of reference and master profiles and their constituent reference characteristics (FIG. 9), and several illustrative update sequences (FIG. 10). With this foundation, reference is made to FIG. 11 to describe the operations of an exemplary update sequence 1100. For ease of explanation, but without any intended limitation, the example of FIG. 11 is described in the context of the hardware environment described above in FIGS. 1-4.

The sequence 1100 further embellishes the operations involved in updating a managed unit's operating characteristics, subscribed-to profiles, and master profiles. In one flow diagram, the sequence 1100 explains the duties of both configuration manager and managed unit involved in an update; thus, the routine 1100 expands previously described operations shown in step 716 (performed by the configuration manager) and steps 816-818 (performed by the managed unit). Accordingly, the routine 1100 starts in step 1102, which occurs when the update criteria are satisfied, as illustrated in step 714 (FIG. 7) and step 814 (FIG. 8).

#### Identifying Stale Subscribed-to Profiles

After step 1102, the managed unit 400 transmits its list of subscribed-to profiles 410 to the configuration manager 300. The reason for sending the subscribed-to profiles 410 is to permit the configuration manager 300 to compare the subscribed-to profiles 410 to the reference profiles 308, and detect any "stale" subscribed-to profiles. "Stale" subscribed-to profiles are defined as any subscribed-to profiles that are older than their counterparts in the list of reference profiles 308. Staleness is detected by comparing version codes of each subscribed-to profile and the reference profile counterpart with the same name. As explained below, the configuration manager detects stale subscribed-to profiles in step 1106, and this fact is used to update the operating characteristics of subscribed-to profiles (step 1108-1112, 1113) and to update the subscribed-to profiles themselves (step 1121).

As an alternative to steps 1104-1106 as illustrated, the configuration manager 300 may send the managed unit 400 its list of reference profiles 308 (step 1104). In this embodiment, the managed unit 400 compares the configuration manager's reference profiles 308 to the locally-stored subscribed-to profiles 410 (step 1106). For explanatory purposes, the current illustration depicts comparison by the configuration manager 300.

Whether performed by the configuration manager or the managed unit, comparison of subscribed-to profiles to reference profiles is performed as follows. Each subscribed-to profile's version code is compared to the version code of its reference profile counterpart to determine whether the subscribed-to profile is current. If the subscribed profile's version code is older than the reference profile's version code, the subscribed-to profile is "stale."

If the comparison of step 1106 shows that a subscribed-to profile's reference profile counterpart no longer exists in the list 308, then the subscribed-to profile is an "orphan" profile. In this event, the managed unit prompts a system administrator for input on how to proceed.

#### Detecting and Implementing New or Updated Operating Characteristics

Having made the necessary comparison to identify any stale subscribed-to profiles, step 1108 asks whether any stale subscribed-to profiles were found. If so, steps 1110-1112

proceed to analyze the differences in each subscribed-to profile to identify any new or updated operating characteristics that must be downloaded to the managed unit. Steps 1110–1112 are performed for each stale subscribed-to profile. Steps 1110–1112 are performed to ensure that the managed unit 400 contains the latest versions of the operating characteristic versions in its list 408.

The detailed comparison of profile contents is performed by comparing the reference profile against the version code of the subscribed-to profile (which was sent to the configuration manager 300 in step 1104). In step 1110, the configuration manager 300 compares the version code of the stale subscribed-to profile to that of its reference profile counterpart. Particularly, step 1110 examines the characteristic version codes 960–964 (FIG. 9) of each reference characteristic 953–957 to determine whether any version code is more recent than the version code of the stale subscribed-to profile. If so, this means the subscribed-to profile refers to a stale operating characteristic version. For each stale operating characteristic, the configuration manager 400 sends an up-to-date version of the counterpart reference characteristic to the managed unit 400 (step 1112). In response, the managed unit 400 stores the received reference characteristic in its list 408 in replacement of the stale operating characteristic.

Step 1110 also identifies any “new” reference characteristics that have been added to the reference profile without addition to the master profile yet (step 1110). For each added reference characteristic, the configuration manager 400 sends the reference characteristic to the managed unit 400 (step 1112).

#### Downloading Profile Master View

After step 1112, the configuration manager 300 downloads all reference profiles to the managed unit 400 (step 1114). This is performed so that (1) the managed unit can compare its profile master view to the up-to-date reference profiles to detect any deleted reference characteristics, and also (2) for use in updating the managed unit's profile master view (as discussed below). The downloaded reference profiles are stored without deleting or changing contents of the existing, but possibly out-of-date, profile master view 409.

#### Detecting and Purging Orphan Operating Characteristics

After step 1114, step 1118 asks whether there are any stale subscribed-to profiles; this was determined in step 1106, discussed above. If not, there is no need to search for deleted reference characteristics, because none of the subscribed-to profiles are stale. In this case, the routine 1100 advances to step 1116, which is discussed below.

If step 1118 finds any stale subscribed-to profiles, then the routine 1100 proceeds to step 1113, which begins the operation of analyzing the master profiles to detect any reference characteristics that have been deleted from the reference profiles 308, but remain in the master profiles 409. The comparison of step 1113 searches for any operating characteristics that are present in the stale master profile but missing from the counterpart reference profile. When the stale master profile is later updated (as described below), this operating characteristic will be deleted from the master profile. Also, if no other reference profiles contain the operating characteristic, it is an “orphan.” In this event, the managed unit step 1113 deletes the orphan operating characteristic from the list 408.

#### Updating Version of Subscribed-to Profile

After step 1113, the managed unit 400 updates the version codes of each stale subscribed-to profile in the list 410 (step 1121). This may be performed, for example, by copying the version code 902 of the corresponding reference profile,

which was sent in step 1114. Also in step 1121, the managed unit reports the new version codes to the configuration manager.

#### Updating the Subscription List

After step 1121, the configuration manager 300 in step 1122 updates its subscription list 312 in accordance with the new version codes of subscribed-to profiles, which were sent in step 1121.

#### Updating Profile Master View

After step 1122 (or a negative answer to step 1118), step 1116 updates the managed unit's profile master view 409 in accordance with the up-to-date reference profiles that were downloaded in step 1114. As an example, step 1116 may be performed by deleting contents of the storage 409, and storing the downloaded reference profiles 308 there. As discussed below, the local presence of the profile master view helps the managed unit more efficiently analyze and update its subscribed-to profiles 410.

After step 1116, the routine 1100 ends in step 1126. At this point, the managed unit returns to step 814 (FIG. 8) of its operating sequence 800, whereas the configuration manager returns to step 712 (FIG. 7).

#### Handling New Subscription Requests

With comparative update tracking as described above, the invention may employ an alternative process for subscription by a managed unit. As described above, a managed unit subscribes to a reference profile by submitting a subscription message in step 810 (FIG. 8). As an alternative, the managed unit may store a subscribed-to profile having a version code of zero in the list 410 without sending any subscription to the configuration manager. This subscribed-to profile identifies the name of a reference profile for which subscription is desired, and lists a version code of zero.

In this embodiment, the configuration manager always uses positive integers when adding new reference profiles. Thus, during the routine 1100, any subscribed-to profiles with version code zero are found (steps 1106, 1108) to be stale. In the manner discussed above, this triggers the configuration manager to send the managed unit all reference characteristics of the subscribed-to profile, which the managed unit proceeds to store in its list 408 (step 1112). The managed unit also updates the corresponding master profile according to the latest version of the reference profile, and stores the new master profile in the list 409 (step 1116). Also, the managed unit updates its subscribed-to profile to bear the up-to-date reference profile's version code (step 1121). Thus, in this embodiment, subscription is incorporated into the existing update process.

#### OTHER EMBODIMENTS

While the foregoing disclosure shows a number of illustrative embodiments of the invention, it will be apparent to those skilled in the art that various changes and modifications can be made herein without departing from the scope of the invention as defined by the appended claims. Furthermore, although elements of the invention may be described or claimed in the singular, the plural is contemplated unless limitation to the singular is explicitly stated.

What is claimed is:

1. In a multiprocessing system that includes:  
a configuration manager maintaining the following data:  
reference data, comprising:  
a record of reference characteristics;  
a record of reference profiles, each reference profile  
having a name and an associated profile version

19

code, each reference profile identifying one or more constituent reference characteristics; and a subscription list identifying each managed unit subscribing to any of the reference profiles, and in association with each subscribing managed unit, all reference profiles subscribed to and their profile version codes;

multiple managed units, each coupled to the configuration manager, each managed unit maintaining the following data:

- a record of one or more operating characteristics;
- a record of one or more subscribed-to profiles, each subscribed-to profile comprising a name of a counterpart reference profile maintained by the configuration manager at a past time, and the counterpart reference profile's profile version code; and
- a profile master view containing master copy counterparts of all reference profiles maintained by the configuration manager at a past time;

where each managed unit functions according to contents of its record of operating characteristics;

a method of comparative update tracking to change local configuration settings at the managed units for synchronization with centrally maintained configuration settings at the configuration manager, comprising operations of:

- the configuration manager receiving changes to the reference data;
- in response to the changes, determining whether profile version codes of any of the subscribed-to profiles are older than profile version codes of their reference profile counterparts currently being maintained by the configuration manager;
- if so, comparing the master copy counterpart of the subscribed-to profile to the reference profile counterpart of the subscribed-to profile as maintained by the configuration manager to identify any differences there between, and
- in response to finding any new reference characteristics identified by the reference profile counterpart of the subscribed-to profile but not identified by the master copy counterpart of the subscribed-to profile, copying the new reference characteristics to the record of operating characteristics;
- in response to finding any updated reference characteristics identified by the reference profile counterpart of the subscribed-to profile where previous versions of the reference characteristic are identified by the master copy counterpart of the subscribed-to-profile, copying the updated reference characteristics to the record of operating characteristics; and
- in response to finding any omitted reference characteristics not identified by the reference profile counterpart of the subscribed-to profile but identified by the master copy counterpart of the subscribed-to-profile, determining whether any other master copy counterparts of subscribed-to profiles refer to the omitted reference characteristic, and if not, deleting the omitted reference characteristic from the record of operating characteristics.

2. The method of claim 1, the method further comprising: if profile version codes of any of the subscribed-to profiles are older than profile version codes of their counterpart reference profile counterparts currently being maintained by the configuration manager,

20

changing the older version code of the subscribed-to profile to match the profile version code of its counterpart reference profile currently being maintained by the configuration manager;

changing the version code of the reference profile as listed in the subscription list in association with the managed unit; and

updating the master copy counterpart of the subscribed-to profile in the profile master view to replicate the reference profile.

3. The method of claim 2, the updating of the counterpart of the subscribed-to profile in the profile master view comprising copying all reference profiles from the record of reference profiles of the configuration manager into the profile master view.

4. The method of claim 1, further comprising, responsive to profile version codes of any of the subscribed-to profiles being older than profile version codes of their counterpart reference profiles currently being maintained by the configuration manager, the managed unit sending contents of its record of subscribed-to profiles to the configuration manager;

where the comparing comprises:

- the configuration manager examining the master copy counterpart of the subscribed-to profile to the reference profile counterpart of the subscribed-to profile as maintained by the configuration manager to identify new reference characteristics listed in the reference profile but not listed in the master copy counterpart; and
- the configuration manager examining the master copy counterpart of the subscribed-to profile to the reference profile counterpart as maintained by the configuration manager to identify updated reference characteristics listed in the reference profile where previous versions of the reference characteristic are listed in the master copy counterpart.

5. The method of claim 4, where each reference profile includes, in association with each constituent reference characteristic, a characteristic version code comprising the profile version code of the reference profile at a time the constituent reference characteristic was added to the reference profile; where the identification of updated reference characteristics comprises comparing the each characteristic version code in the reference profile counterpart to the profile version code of the master copy counterpart of the subscribed-to profile.

6. The method of claim 1, further comprising duplicating the reference profiles from the configuration manager at the managed unit; where the comparing includes the managed unit comparing the master copy counterpart of the subscribed-to profile to the duplicate reference profile counterpart of the subscribed-to profile to identify any omitted reference characteristics not identified by the master copy counterpart of the subscribed-to-profile, and also determining whether any other master copy counterparts of subscribed-to profiles refer to the omitted reference characteristic, and if not, deleting the omitted reference characteristic from the record of operating characteristics.

7. The method of claim 6, further comprising: replacing contents of the profile master view with the reference profiles duplicated from the configuration manager.

21

8. The method of claim 1, each profile version code comprising an integer, the method further including incrementing a reference profile's version code each time the reference profile is changed.

9. The method of claim 1, further comprising:  
the managed unit receiving a subscription request identifying a desired reference profile;  
in response to the subscription request, the managed unit representing the desired reference profile in the record of subscribed-to profiles.

10. The method of claim 9, where the desired reference profile has a name, and where the managed unit representing the desired reference profile in the record of subscribed-to profiles comprises storing the name in the record of subscribed-to profiles in association with a predetermined version code earlier than all previous versions of the desired reference profile.

11. The method of claim 10, the predetermined version code being zero.

12. The method of claim 1, further comprising:  
the managed unit receiving operating characteristics from source other than the configuration manager, and in response, storing the received operating characteristics in the record of operating characteristics.

13. A multiprocessing system, comprising:  
a configuration manager maintaining the following data:

reference data, comprising:  
a record of reference characteristics;  
a record of reference profiles, each reference profile having a name and an associated profile version code, each reference profile identifying one or more constituent reference characteristics; and  
a subscription list identifying each of the managed units subscribing to any of the reference profiles, and in association with each subscribing managed unit, all reference profiles subscribed to and their profile version codes;

one or more managed units, each coupled to the configuration manager, each managed unit maintaining the following data:

a record of one or more operating characteristics;  
a record of one or more subscribed-to profiles, each subscribed-to profile comprising a name of a counterpart reference profile maintained by the configuration manager at a past time, and the counterpart reference profile's profile version code; and  
a profile master view containing master copy counterparts of all reference profiles maintained by the configuration manager at a past time;

where each managed unit functions according to contents of its record of operating characteristics;

a communications link coupling the configuration manager to the managed units;

where the configuration manager and managed units are programmed to perform operations of comparative update tracking to synchronize local configuration settings with centrally maintained paradigm settings, the operations comprising:

the configuration manager receiving changes to the reference data;

in response to the changes, determining whether profile version codes of any of the subscribed-to profiles are older than profile version codes of their reference profile counterparts currently being maintained by the configuration manager;

if so, comparing the master copy counterpart of the subscribed-to profile to the reference profile coun-

22

terpart of the subscribed-to profile as maintained by the configuration manager to identify any differences there between, and  
in response to finding any new reference characteristics identified by the reference profile counterpart of the subscribed-to profile but not identified by the master copy counterpart of the subscribed-to profile, copying the new reference characteristics to the record of operating characteristics;

in response to finding any updated reference characteristics identified by the reference profile counterpart of the subscribed-to profile where previous versions of the reference characteristic are identified by the master copy counterpart of the subscribed-to-profile, copying the updated reference characteristics to the record of operating characteristics; and  
in response to finding any omitted reference characteristics not identified by the reference profile counterpart of the subscribed-to profile but identified by the master copy counterpart of the subscribed-to-profile, determining whether any other master copy counterparts of subscribed-to profiles refer to the omitted reference characteristic, and if not, deleting the omitted reference characteristic from the record of operating characteristics.

14. The system of claim 13, the configuration manager and managed units being further programmed to perform operations comprising:

if profile version codes of any of the subscribed-to profiles are older than profile version codes of their counterpart reference profile counterparts currently being maintained by the configuration manager,  
changing the older version code of the subscribed-to profile to match the profile version code of its counterpart reference profile currently being maintained by the configuration manager;  
changing the version code of the reference profile as listed in the subscription list in association with the managed unit; and  
updating the master copy counterpart of the subscribed-to profile in the profile master view to replicate the reference profile.

15. The system of claim 14, the configuration manager and managed units being programmed such that the updating of the counterpart of the subscribed-to profile in the profile master view comprises copying all reference profiles from the record of reference profiles of the configuration manager into the profile master view.

16. The system of claim 13,  
the configuration manager and managed units being further programmed to perform operations further comprising, responsive to profile version codes of any of the subscribed-to profiles being older than profile version codes of their counterpart reference profiles currently being maintained by the configuration manager, the managed unit sending contents of its record of subscribed-to profiles to the configuration manager;

where the configuration manager and managed units are programmed such that the comparing comprises:  
the configuration manager examining the master copy counterpart of the subscribed-to profile to the reference profile counterpart of the subscribed-to profile as maintained by the configuration manager to identify new reference characteristics listed in the reference profile but not listed in the master copy counterpart; and

UNITED STATES PATENT AND TRADEMARK OFFICE  
CERTIFICATE OF CORRECTION

PATENT NO. : 6,167,408  
DATED : December 26, 2000  
INVENTOR(S) : Cannon et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Title page.

Change "UPDATES" to -- UPDATE --

Column 2.

Line 46, change "subscribed-o" to -- subscribed-to --

Column 7.

Table 2, approx. line 15, change "Address=9.115.12.92 Port=1500" to  
-- Address=9.115.12.92, Port=1500 --

Column 12.

Lines 24-25, between "input," and "rather" delete "5"

Claim 6, column 20.

Line 58, change "subscribed-to-profile" to -- subscribe-to profile --

Claim 13, column 22.

Line 14, change "subscribed-to-profile" to -- subscribe-to profile --

Claim 13, column 22.

Line 21, change "subscribed-to-profile" to -- subscribe-to profile --

Claim 18, column 23.

Line 32, change "subscribed-to-profile" to -- subscribe-to profile --

Signed and Sealed this

Thirtieth Day of October, 2001

Attest:

*Nicholas P. Godici*

Attesting Officer

NICHOLAS P. GODICI  
Acting Director of the United States Patent and Trademark Office



US006266690B1

(12) **United States Patent**  
Shankarappa et al.

(10) Patent No.: **US 6,266,690 B1**  
(45) Date of Patent: **Jul. 24, 2001**

(54) **ENHANCED SERVICE PLATFORM WITH SECURE SYSTEM AND METHOD FOR SUBSCRIBER PROFILE CUSTOMIZATION**

(75) Inventors: Vijay Shankarappa, San Jose, CA (US); Eric William Burger, McLean, VA (US); John Nestoriak, III, Bethesda, MD (US)

(73) Assignee: ADC Telecommunications, Inc., Minnetonka, MN (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: 09/361,676

(22) Filed: **Jul. 27, 1999**

**Related U.S. Application Data**

(63) Continuation-in-part of application No. 09/238,108, filed on Jan. 27, 1999.

(51) Int. Cl.<sup>7</sup> ..... **G06F 15/16**

(52) U.S. Cl. ..... **709/202; 709/202; 709/203; 709/227; 709/229; 707/1; 707/10; 707/100; 707/200; 455/461; 370/351; 370/385**

(58) **Field of Search** ..... **709/202-203; 707/1, 10, 100, 200; 455/461; 370/351, 385**

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

5,610,910 * 3/1997 Focsanecanu et al. ....	370/351
5,742,905 * 4/1998 Pepe et al. ....	455/461
5,774,671 * 6/1998 Satoh ....	709/231

5,813,006 * 9/1998 Polnerow et al. ....	707/10
5,867,498 * 2/1999 Gillman et al. ....	370/385
5,881,234 * 3/1999 Schwob ....	709/219
5,898,839 * 4/1999 Berreau ....	709/227
5,958,016 * 9/1999 Chang et al. ....	709/229

**OTHER PUBLICATIONS**

PCT/ISA International Search Report, Nov. 24, 2000, International Application No. PCT/US00/20434, 4 Sheets.

\* cited by examiner

*Primary Examiner*—Ayaz Sheikh

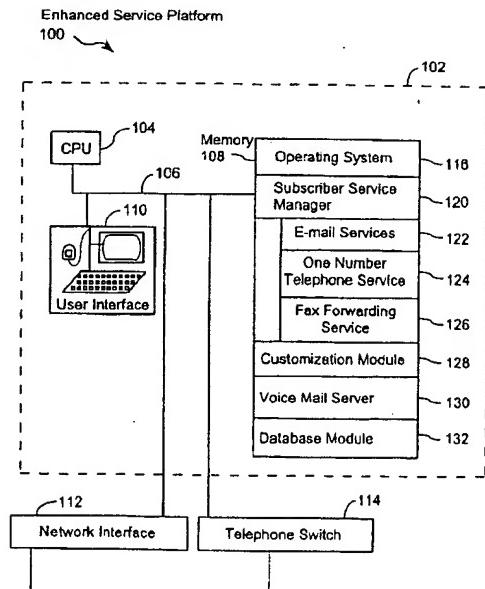
*Assistant Examiner*—Thu Ha Nguyen

(74) *Attorney, Agent, or Firm*—Pennie & Edmonds LLP

(57) **ABSTRACT**

An enhanced telephone service system including a subscriber service manager and a customization module. The subscriber service manager is configured to receive subscriber transaction-ID requests, supply transaction identifiers and respond to queries for subscriber profile data. The customization module is configured to receive customization messages. When a customization message includes a transaction identifier, the customization module compares the transaction identifier in the customization message with the transaction identifiers generated by the subscriber service manager. Only if there is a match between the transaction identifiers does the customization module register the customization data. Alternatively, if the customization message does not contain a transaction identifier but does contain a subscriber identifier, the customization module temporarily stores the customization data. When a subscriber logs on to the system and there is pending temporarily stored customization data, the customization module prompts the subscriber for authorization to register the temporarily stored data.

33 Claims, 9 Drawing Sheets



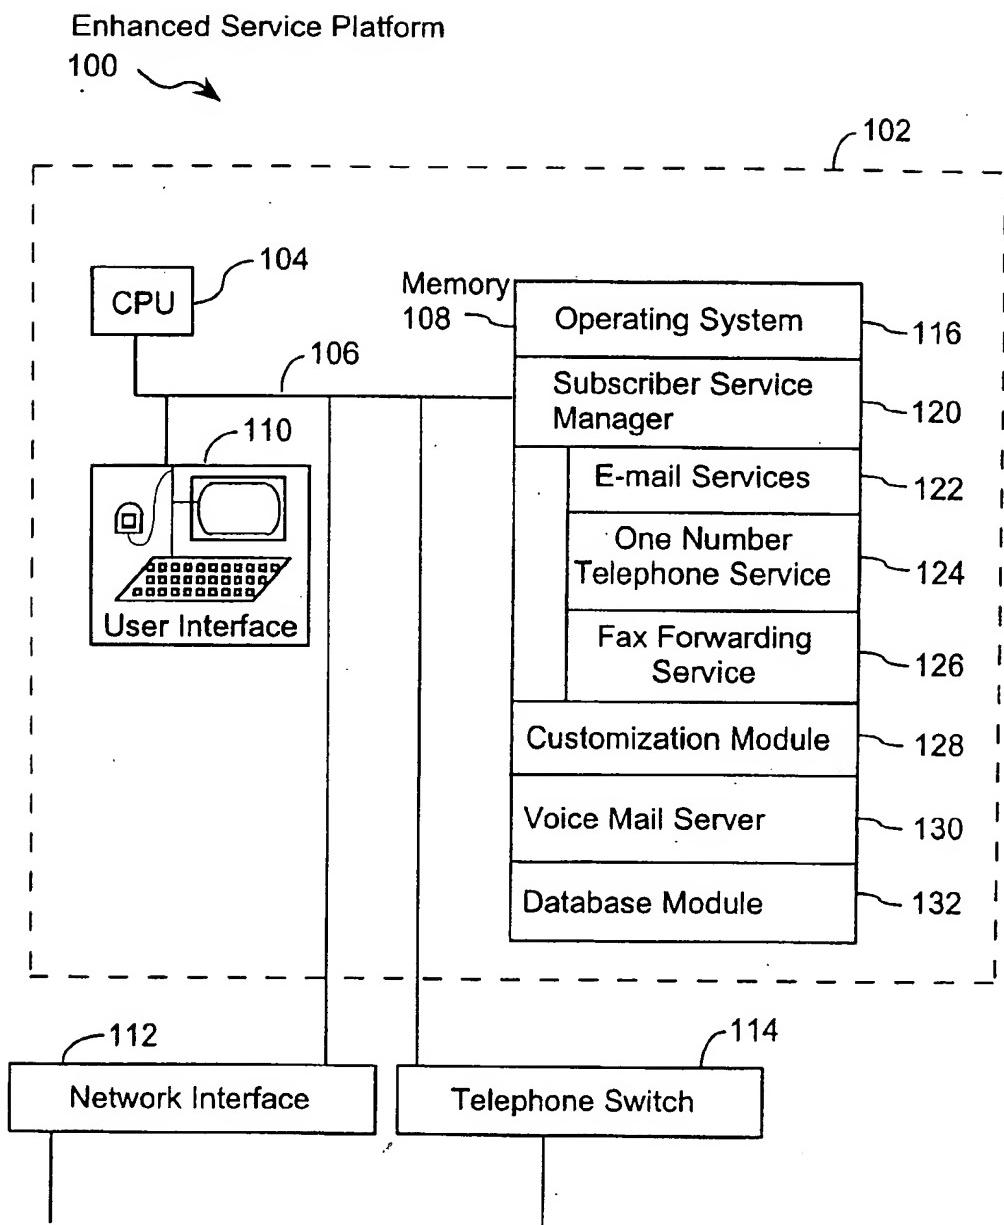


FIG. 1

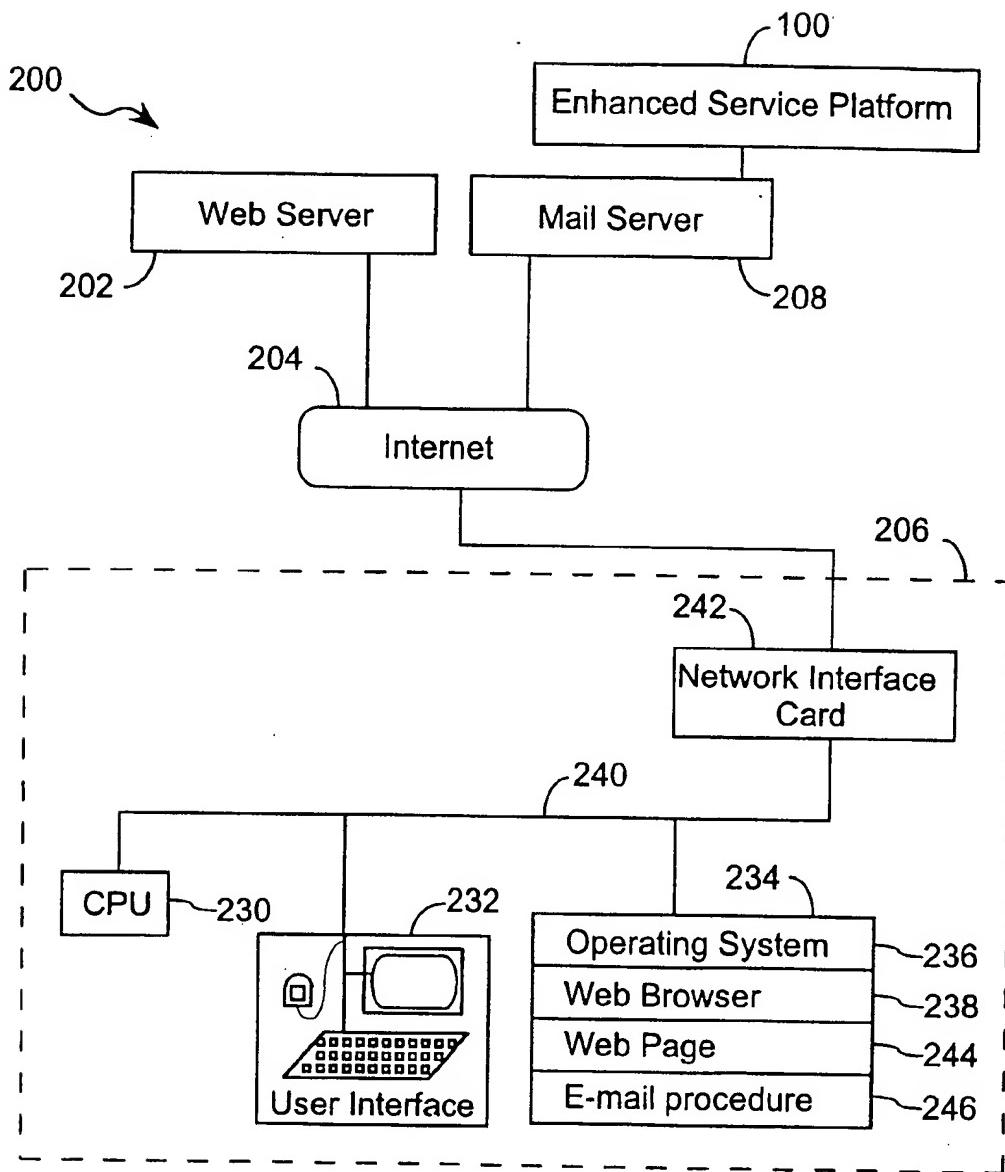


FIG. 2

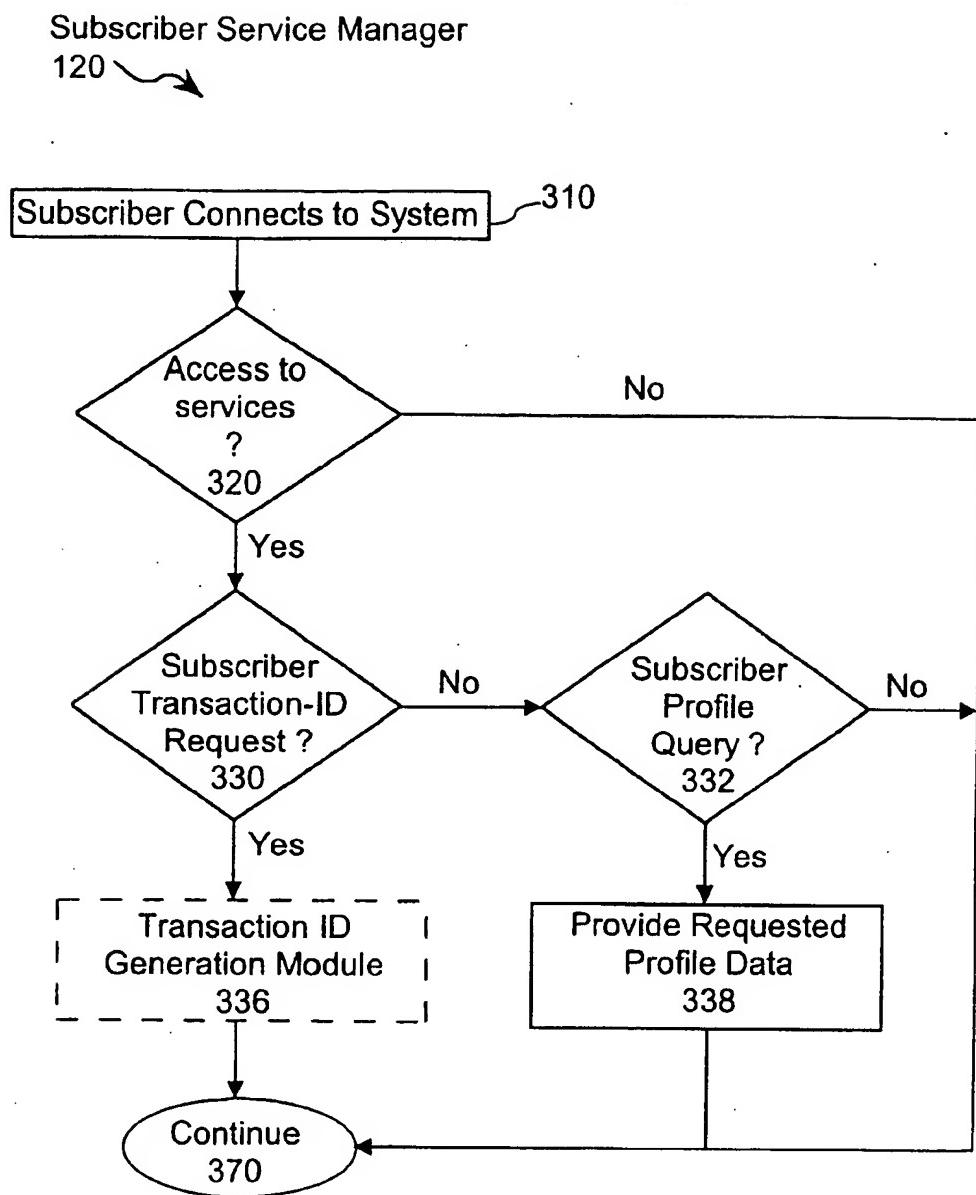


FIG. 3A

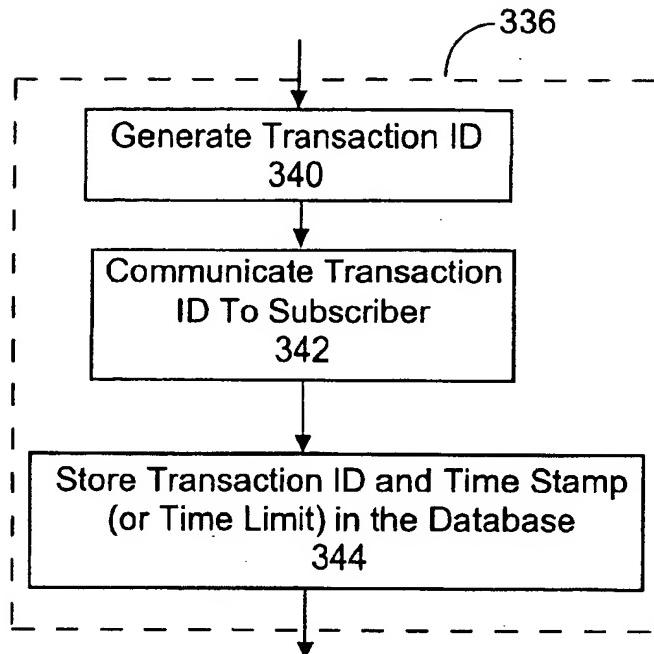


FIG. 3B

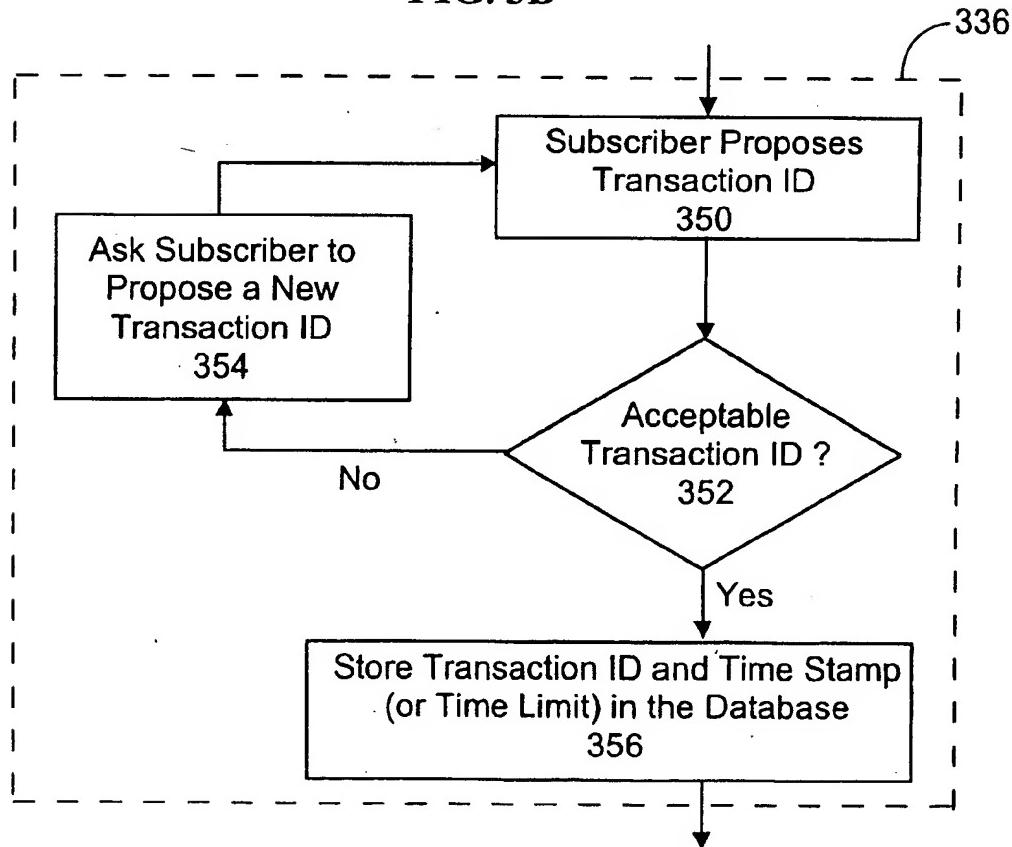


FIG. 3C

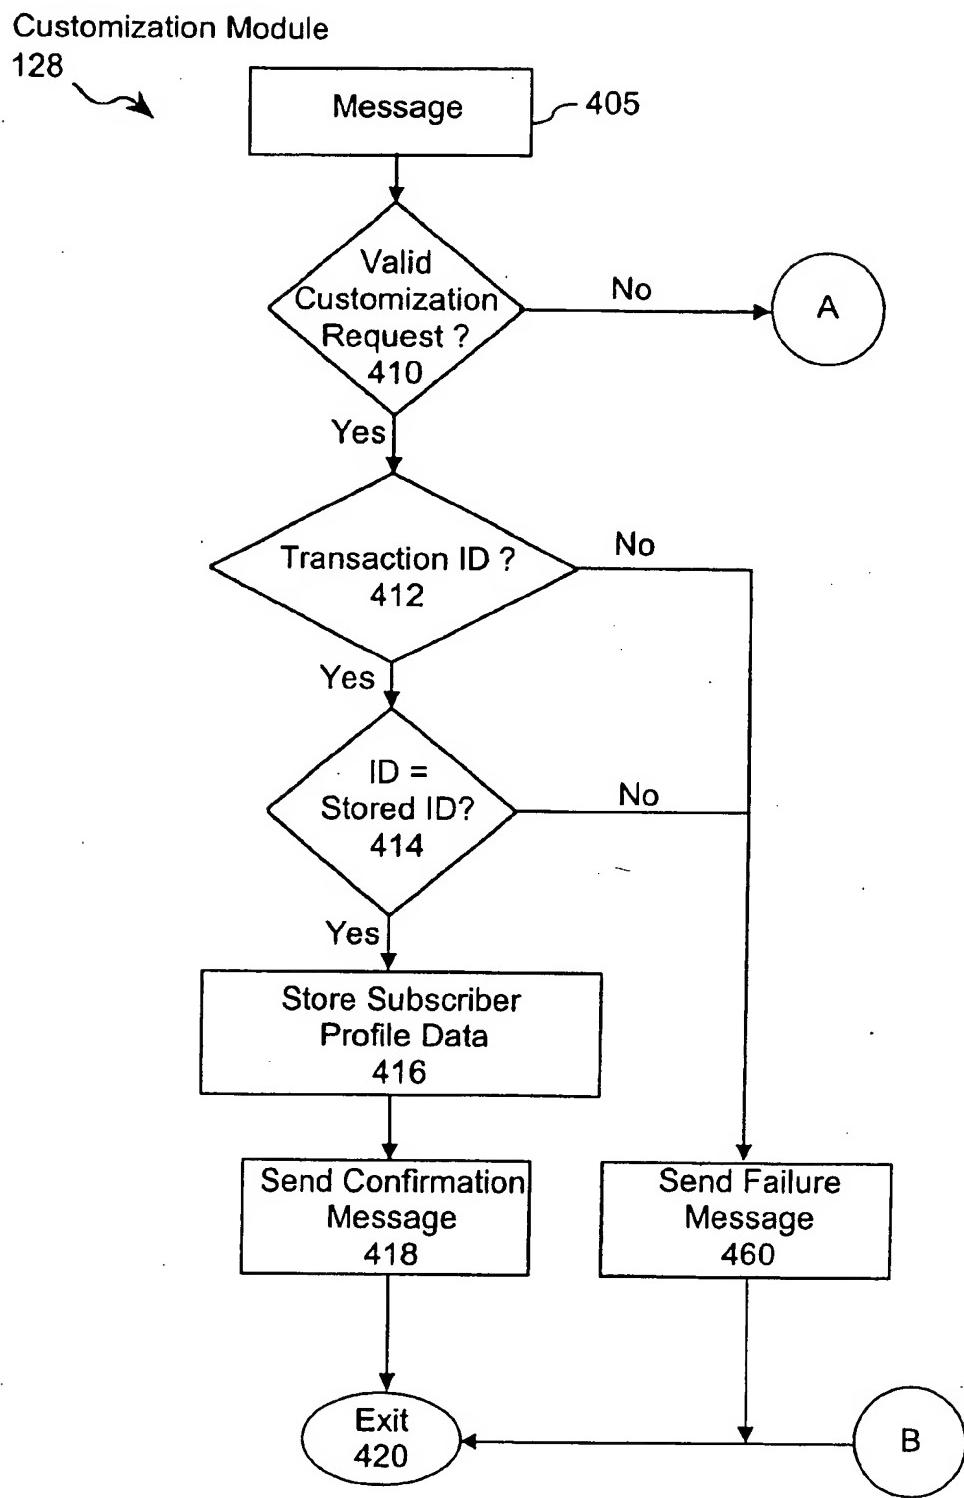


FIG. 4A

## Customization Module (continued)

128

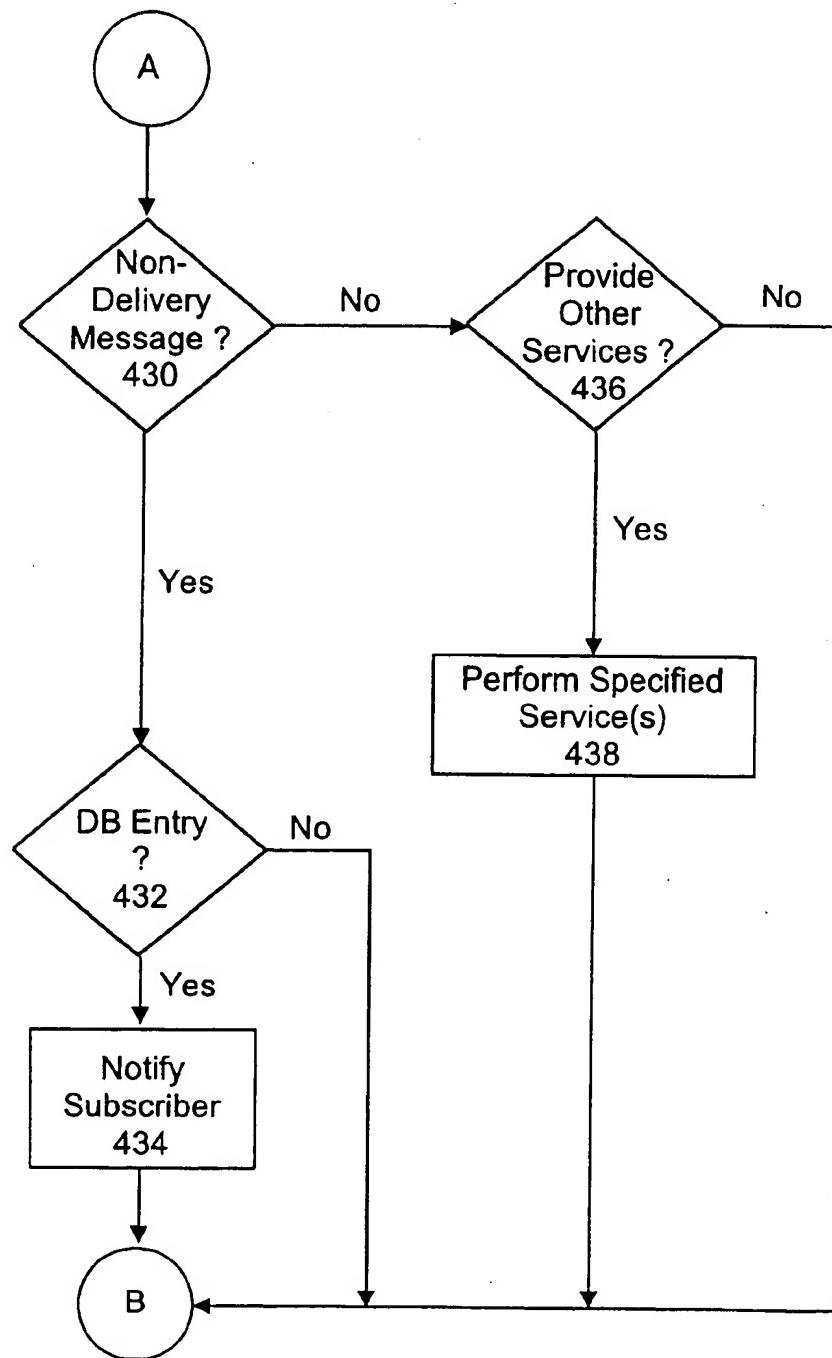


FIG. 4B

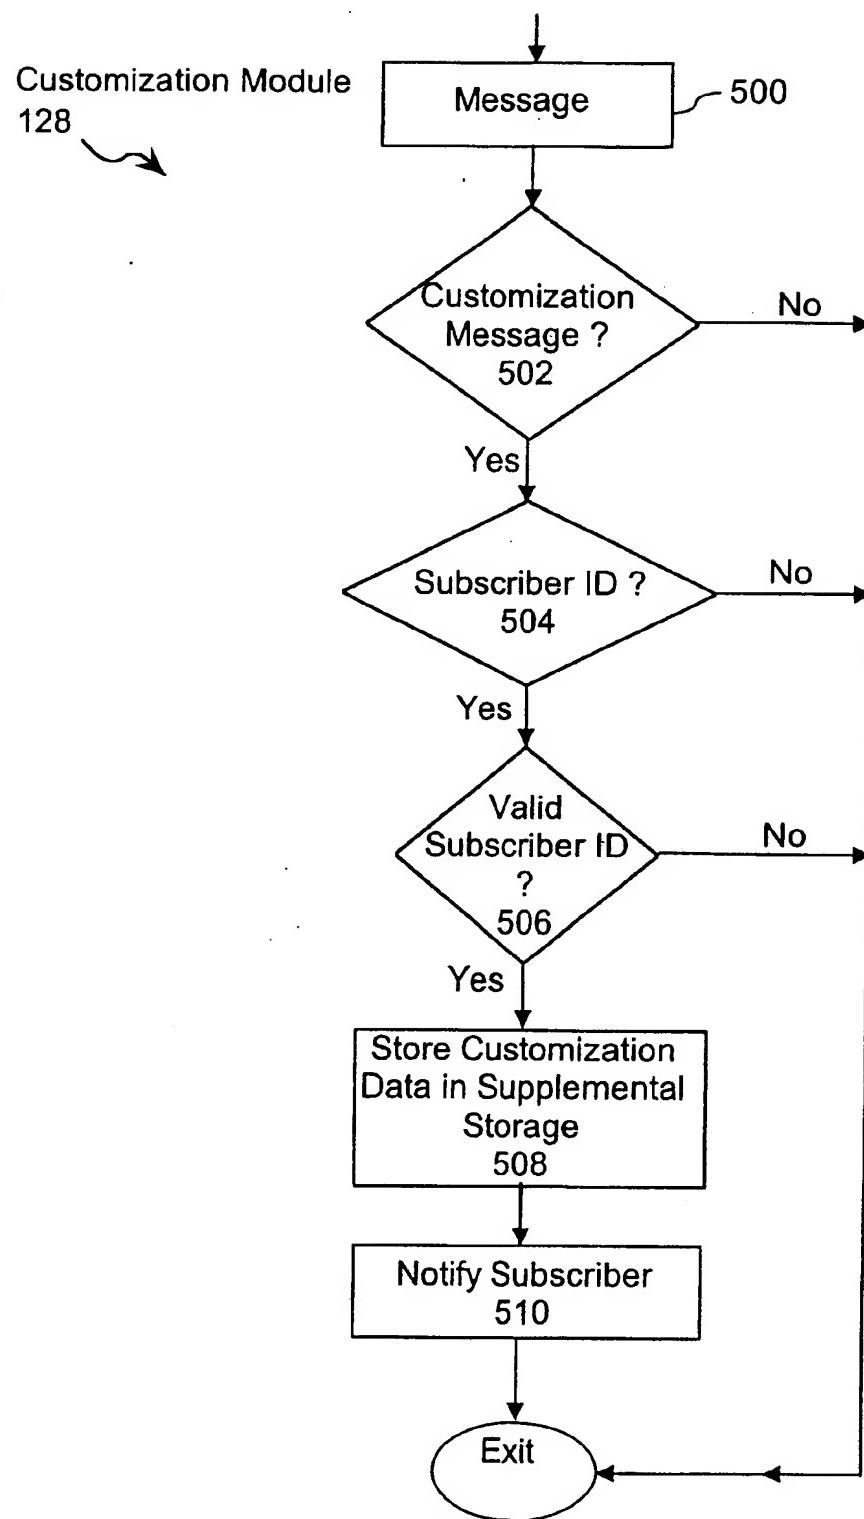


FIG. 5A

Customization Module  
128

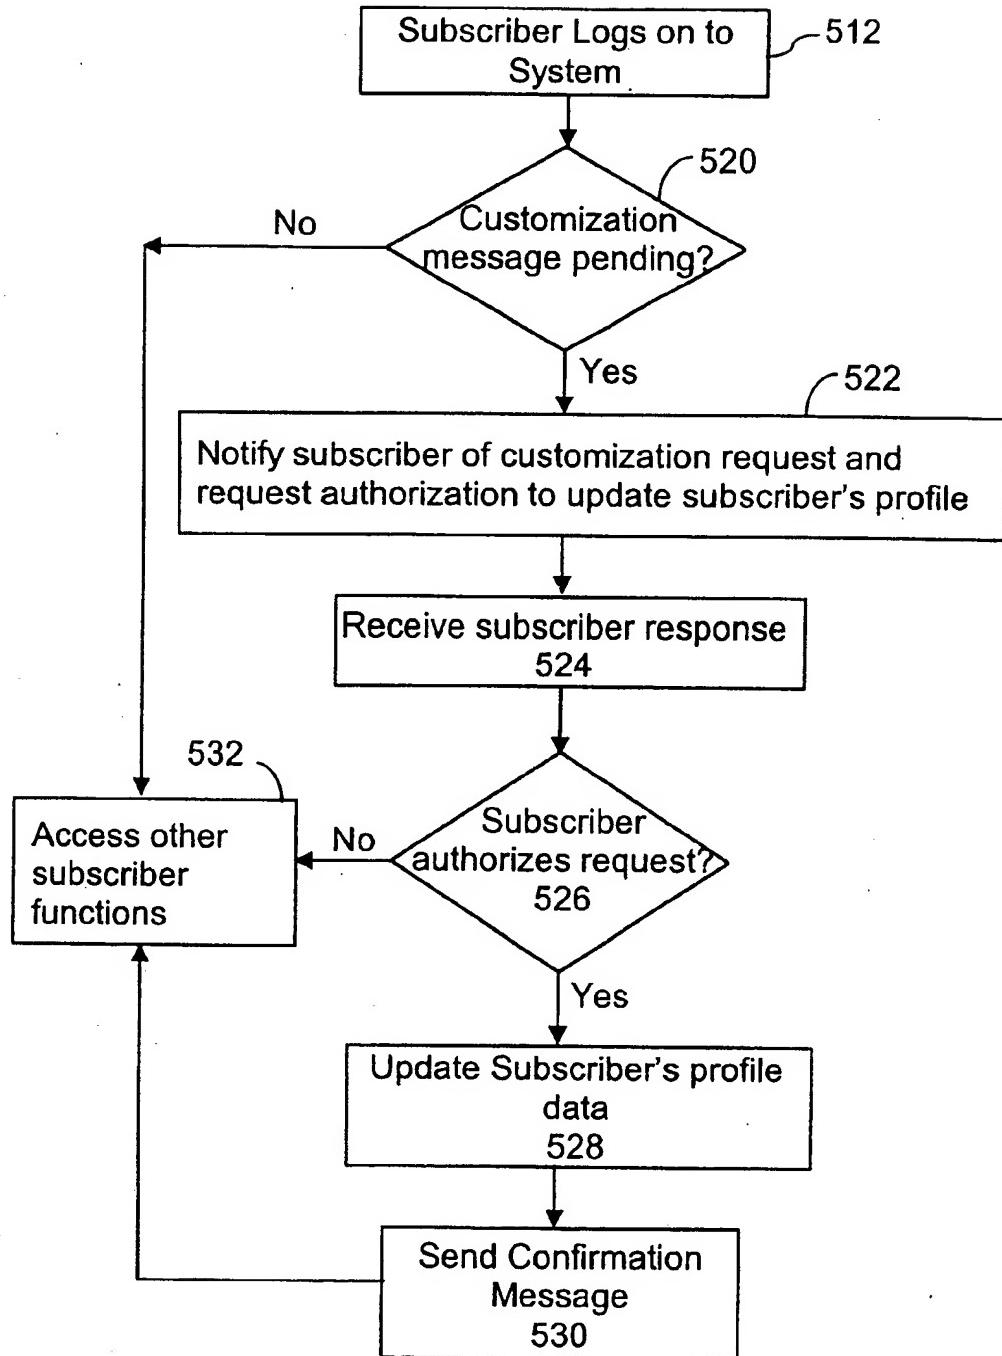


FIG. 5B

Database Record

600

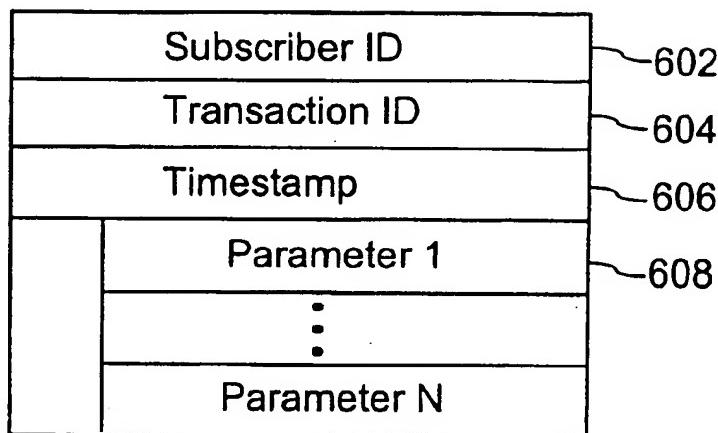


FIG. 6A

Database Record

650

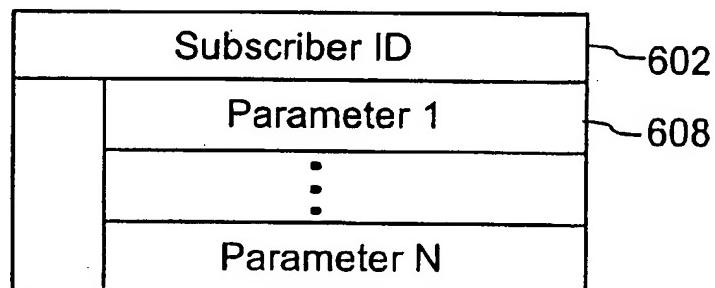


FIG. 6B

**ENHANCED SERVICE PLATFORM WITH  
SECURE SYSTEM AND METHOD FOR  
SUBSCRIBER PROFILE CUSTOMIZATION**

This application is a continuation-in-part of U.S. patent application Ser. No. 09/238,108, filed Jan. 27, 1999, entitled "Enhanced Telephone Service System With Secure System And Method For E-mail Address Registration."

The present invention relates generally to a subscriber profile customization system and method and more particularly to a secure method for registering subscriber profile data with an enhanced service provider.

**BACKGROUND OF THE INVENTION**

Registration of subscriber profile data is well known among service providers. Service providers routinely require subscribers to specify subscriber profile data such as a personal E-mail (electronic mail) address. This is typically completed by having the subscriber do one of the following: send an E-mail message containing subscriber profile data to a specified E-mail address; use a web browser, such as Netscape Navigator, to log onto the service provider's web site on the world-wide-web (WWW) to register the data; spell out the data on a telephone keypad; or recite the data to a live operator.

A list server is an example of a service that often requires subscribers to register subscriber profile data such as a personal E-mail address in order to receive mail deliveries. A subscriber is generally required to send an E-mail message to a specified E-mail address, the E-mail message requesting subscription to the list server's services. After receiving the E-mail message, the list-server registers the corresponding E-mail address by storing the network routing information contained in the E-mail message. For more information on list-servers, see *List Owner's Manual for LISTSERV, ver1.8*, L-Soft International, Inc., 1997, which is hereby incorporated by reference for purposes of background information.

The problems associated with known subscriber profile data registration techniques are numerous. For example, notifying a subscriber that the registration attempt was, or was not, successful is often difficult; entering subscriber profile data such as an E-mail address via a telephone keypad is often tedious and error prone; and, verifying that a subscribed service is being directed to a legitimate entity is difficult.

After processing a subscriber profile registration request from a subscriber, notifying the subscriber that the registration request was, or was not, successful can be difficult. A common method of registration verification is for a service provider to send an E-mail message to a subscriber's E-mail address confirming the registration. The network routing information contained in the subscribing profile registration request is typically used for this purpose. However, lack of a corresponding non-delivery notification does not positively indicate successful delivery of the confirming E-mail message. To further complicate matters, if the E-mail address is invalid for any reason, a service provider typically has no way of notifying the subscriber to try to register again.

Using a telephone keypad to enter subscriber profile data is not a very accurate method of registering subscriber data. Subscriber profile data such as E-mail addresses are often quite long and typically include non-alphanumeric characters, such as the "at" sign (@) or the underscore (\_); neither punctuation is represented on a telephone keypad. Therefore, the likelihood of error while typing in subscriber

profile data on a telephone key pad is quite high. Similarly, speaking subscriber profile data to a live telephone operator or to a voice-recognition system is also error prone.

It is important to ensure that a subscribed service is being used by legitimate subscribers and similarly that the subscriber's profile customization data is not being altered by nonsubscribers. This is especially pertinent when the services being provided are proprietary or expensive. Conventional enhanced service platforms typically use an external web server to communicate with subscribers. The external web server typically communicates with subscribers via a direct Internet connection and with the enhanced service provider via an Intranet connection. While a direct Internet connection between subscribers and the external web server provides a convenient data exchange mechanism, the direct Internet connection is vulnerable to third party attack. Consequently, both the Internet and the Intranet connections must be regulated by expensive firewall devices in order to insure the security of the data transmitted to the enhanced service platform. These firewall devices are expensive, limit the number and types of communication protocols that may be used by the web server, and remain sensitive to attack by third parties.

What is needed is a subscriber profile registration system and method for registering subscriber profile data with a service provider without the data specification, verification, and notification issues commonly encountered by service provider platforms. Thus, a system is needed that provides a method for accurately entering subscriber profile data, provides a method to verify that subscriber profile data is being modified by a legitimate entity, and insures that the status of a successful or unsuccessful subscriber profile registration attempt is conveyed to the subscriber. Preferably, the system will also allow a subscriber to query subscriber profile data.

**BRIEF DESCRIPTION OF THE DRAWINGS**

Additional objects and features of the invention will be more readily apparent from the following detailed description and appended claims when taken in conjunction with the drawings, in which:

FIG. 1 is a block diagram of an embodiment of the enhanced service platform of the present invention, for subscriber profile registration with a service provider;

FIG. 2 is a block diagram of a subscriber workstation with web and mail server interfaces for entering subscriber profile data that is to be registered with a enhanced service platform or query registered profile data;

FIG. 3A is a flow diagram of an embodiment of a Subscriber Service Manager for managing subscriber requests at a service provider,

FIGS. 3B and 3C depict alternative embodiments for generating transaction IDs;

FIGS. 4A and 4B depict flow diagrams of a first embodiment of a Customization Module for processing a customization message;

FIGS. 5A and 5B depict flow diagrams of alternative embodiments of a Customization Module for processing a customization message; and

FIG. 6A and 6B depict block diagrams of alternative embodiments of a database record, for storing subscriber profile data that is to be registered with the platform.

**SUMMARY OF THE INVENTION**

The present invention is used in conjunction with a system such as an Enhanced Service Platform 100 that

maintains a profile for each subscriber of services provided by the platform. The system and method of the present invention provides a Subscriber Service Manager and a Customization Module for registering new subscriber profile data and/or customizing preexisting subscriber profile data. Customization includes selecting or deselecting subscriber services offered by the Enhanced Service Platform, or changing data entries in the subscriber profile. The information in the profile may include E-mail addresses, billing information, the names of the services the subscriber subscribes to, personal telephone numbers, and related data that is specific to the subscriber. Subscribers register new data and update previously stored data in their profile using the system and methods of the present invention. Further, subscribers may query their subscriber profile.

The Subscriber Service Manager is configured to receive various requests from the subscriber. In particular, the Subscriber Service Manager is configured to receive a request, termed a "transaction-ID request," to obtain permission to register and/or customize subscriber profile data. Upon receiving such a transaction-ID request, the Subscriber Service Manager supplies the subscriber with a transaction identifier, and stores a copy of the transaction identifier.

The Subscriber Service Manager is also configured to receive subscriber profile query requests. Upon receiving such a request, the Subscriber Service Manager supplies the subscriber with data from the subscriber's profile. Typically the query will be for the data contained in specific fields of the subscriber's profile.

The Customization Module is configured to receive a customization message. The customization message is a message originated by the subscriber that contains new profile data for registration or profile customization. The customization message may be an E-mail or a message generated by web based applications. When a valid customization message is received by the system, the message will specify the subscriber profile data to be registered, as well as the transaction identifier supplied by the Subscriber Service Manager. The Customization Module compares the transaction identifier in the customization message with the copy of the transaction identifier stored by the Subscriber Service Manager. Only if they match does the Customization Module register the subscriber profile data.

In a preferred embodiment, a subscriber first telephones the service platform and provides information that is used to verify that the subscriber has a privileged status that entitles the subscriber to use the platform's services. The subscriber then makes a transaction-ID request to register profile data such as an E-mail address. In response to the transaction-ID request, the service platform supplies a transaction identifier to the subscriber, and stores the transaction identifier, a time-stamp and the subscriber's telephone number for later retrieval. The subscriber then sends a customization message to the service platform requesting that the subscriber's profile be customized or updated with specified subscriber profile data. In addition to subscriber profile data, the customization message includes the transaction identifier. Upon receiving the subscriber's customization message, the Customization Module uses the transaction identifier to verify the customization request. If the customization request is verified, the Customization Module registers the subscriber profile data and/or customizes the subscriber's profile. Then, the Customization Module sends a registration notification to the subscriber, informing the subscriber that new data has been registered in the subscriber's profile and/or the subscriber's profile has been customized.

Preferably, the registration notification is sent to the subscriber's E-mail address, which is registered in the

subscriber's profile. If the registration notification is not deliverable, the Customization Module will attempt to re-deliver the registration notification. If a certain number of such delivery attempts fail, the service platform contacts the subscriber via a telephone number stored in the subscriber's profile to deliver a voice message indicating that the subscriber's registered E-mail address may be invalid or inoperative. In this manner, the service platform solves personal E-mail address specification, verification, and notification issues commonly encountered in subscriber profile registration processes.

#### DESCRIPTION OF THE PREFERRED EMBODIMENTS

##### The Enhanced Service Platform

Referring to FIG. 1, there is shown an embodiment of the Enhanced Service Platform 100 of the present invention. The Enhanced Service Platform 100 may be implemented using a programmed general-purpose computer system 102. The Enhanced Service Platform 100 may include:

- 15 one or more data processing units (CPUs) 104;
- memory 108, which will typically include both high speed random access memory as well as non-volatile memory (such as one or more magnetic disk drives);
- 20 a network or other communication interface 112, for communicating with other computers and other devices;
- a telephone interface 114, for receiving and forwarding telephone calls; and
- 25 one or more communication busses 106, for interconnecting the CPU(s) 104, memory 108, user interface 110, network interface 112, and telephone switch 114.

The computer system's memory 108 stores procedures and data, typically including:

- 30 an operating system 116 for providing basic system services;
- a subscriber service manager 120, for providing telephony services;
- 35 a customization module 128, for providing subscriber data registration and profile customization services;
- 40 a voice mail server 130, for receiving and forwarding voice messages; and
- a database module 132, for storing subscriber profile data.

The subscriber service manager 120 may include:

- 45 an E-mail services module 122, for providing E-mail services to subscribers;
- a one number telephone service 124, for providing a plurality of one-number telephony services to subscribers; and
- 50 a fax forwarding service 126, for forwarding faxes to subscribers via E-mail.

##### Subscriber Workstation and Interfaces to Enhanced Service Platform

Referring to FIG. 2, there is shown a system 200 in which a subscriber workstation 206 communicates with an Enhanced Service Platform 100 via communications network 204 (such as the Internet) and a web server 202 or a mail server 208. In the preferred embodiment, only the mail server 208 is used to direct communications from subscriber workstations 206 to the enhanced service platform 100. The mail server 208 is preferably implemented as a separate server from the enhanced service platform 100, which helps to protect the enhanced service platform 100 from security attacks and avoids the need for an expensive firewall. In an alternate embodiment, the mail server 208 may be integrated with the enhanced service platform 100, in which case additional security equipment may be needed to prevent unwanted communications between other computers and the

enhanced service platform 100. In yet other embodiments, no mail server 208 is provided and enhanced service provider 100 is directly coupled to the communications network 204.

In embodiments that include the web server 202, the web server 202 can be used to generate requests that are routed to the Subscriber Service Manager 120 and customization messages that are routed to the Customization Module 128.

The subscriber workstation 206 may include: one or more data processing units (CPUs) 230; a user interface 232, for inputting data; memory 234, which will typically include both high speed random access memory as well as non-volatile memory (such as one or more magnetic disk drives); a network interface card 242 for connecting communication network 204 to computer system 206; and a bus 240 for interconnecting CPU 230, user interface 232, memory 234, and network interface card 242.

The subscriber workstation's memory 234 stores procedures and data, typically including: an operating system 236 for providing basic system services; a web browser 238 for providing a user interface for accessing communication network 204 and web server 202; a web page 244 that provides options to the subscriber including generation of transaction-ID requests, subscriber profile query requests, and customization messages; and an E-mail procedure 246 for sending and receiving E-mail messages.

In a preferred embodiment, web server 202 is implemented on a different computer platform from the enhanced service platform 100, although in alternate embodiments the web server 202 could be integrated with the enhanced service platform 100. The primary role of web server 202 is to help subscribers compose properly formatted subscriber requests. These requests are sent, by E-mail or other form of electronic communication, to enhanced service provider 100 for processing. For instance, web server 202 may provide subscriber workstations 206 with an HTML form that collects data from the user. This data is submitted to web server 202 as a CGI script, which is then executed by web server 202 to generate an E-mail message that contains the subscriber's request. The E-mail message is sent to Mail server 208. Mail server 208 screens out unauthorized communications and passes authorized E-mail messages from subscribers to enhanced service provider 100. Because E-mail messages are used to communicate with the enhanced service provider 100 rather than a direct Internet connection, in the embodiment shown in FIG. 2, the need for an expensive fire wall is obviated. Specific embodiments of mail server 208 include, but are not limited to, SMTP on a private network, UUCP on a private network, and UUCP on a dial-up network.

One of skill in the art will appreciate that the web page sent by web server 202 to subscriber workstations 206 may contain a Java applet or script that collects the data from the subscriber and composes an E-mail message that is sent from the subscriber's workstation to mail server 208.

#### Subscriber Service Manager

Referring to FIG. 3A, there is shown a flow diagram of an exemplary embodiment of Subscriber Service Manager 120. A subscriber communicates electronically 310 with Subscriber Service Manager 120. This communication 310 may be in the form of a telephone call, E-mail message, or a network connection. Either the subscriber or the Subscriber

Service Manager 120 may originate connection 310. However, in a typical embodiment, the subscriber "logs on" to the Enhanced Service Provider 100 and thereby establishes the electronic communication with the Subscriber Service Manager.

The subscriber is first verified 320 to ensure that the subscriber has legitimate access to Subscriber Service Manager 120 services. When the subscriber does not have legitimate access to such services (320-No), the process may perform other steps 370 (e.g., logging information concerning the unsuccessful attempt and/or establishing a new account) before exiting. If the subscriber does have legitimate access (320-Yes), subscriber service manager 120 is placed in a state in which it will accept transaction-ID requests 330 from the subscriber. If the subscriber is not making a transaction-ID request (330-No), it is possible that the subscriber wishes to query his or her subscriber profile. If so (332-Yes), the Subscriber Service Manager provides specific fields of data that are requested by the subscriber 338. When the subscriber does elect to make a transaction-ID request (330-Yes), the transaction ID generation module 336 performs a number of steps to generate a transaction ID. If the subscriber is neither making a transaction-ID request (330-No) nor making a subscriber profile query (332-No), procedure 370 is called.

There are several alternative embodiments for how a transaction-ID request may be initiated (330) and a transaction ID generated (336). Referring to FIGS. 3A and 3B, if the subscriber has connected to Enhanced Service Platform 100 via a host web page, the option to make a transaction-ID request 330 may be presented to the subscriber as a graphic icon or menu option on a web page. Upon making a transaction-ID request (330-Yes), the Subscriber Service Manager 120 generates a transaction ID 340 (FIG. 3B) and communicates the transaction ID to the subscriber 342 by posting the transaction ID in a reply web page or E-mailing the transaction ID to a secure E-mail account associated with the subscriber. The Subscriber Service Manager 120 then stores the transaction ID and a time stamp in a database 344. The time stamp indicates when the transaction ID was created and is used to determine an expiration date for the transaction ID. Alternatively, a time limit value is stored with the transaction ID.

If the subscriber has connected to Enhanced Service Platform 100 via telephone, a similar process, to that described above, is used to obtain a transaction ID. In a preferred embodiment, options are presented to the subscriber in the form of voice directions, possibly including a predefined hierarchical menu for the subscriber to navigate. The subscriber may select options either by pressing keys on the telephone keypad, or by speaking an appropriate command or response. If Subscriber Service Manager 120 includes a speech recognition driven interface. In the latter case, the system might ask an open question such as "What do you want to do?" The subscriber, in response, might say, "Make a transaction-ID request". Once the transaction-ID request is made (330-Yes), the Subscriber Service Manager generates a transaction ID 340 and recites it to the subscriber 342. Preferably, the transaction ID is recited twice, and the subscriber is given the option of having it recited additional times. Finally, the Subscriber Service Manager stores the transaction ID and a time stamp or a time limit value in a database 344.

In yet another embodiment, if the subscriber has contacted the system via E-mail, the transaction-ID request 330 may be encoded within the header or body of the E-mail message. In such an embodiment, communication of the

Module 128 to provide the disclosed process in a variety of ways. For example, step 412, for determining if message 405 contains a transaction ID, could be performed outside of Customization Module 128.

Referring to FIGS. 5A-5B, there are shown flow diagrams of a preferred embodiment of Customization Module 128. In this preferred embodiment, an alternative procedure can be used for registering subscriber profile data. When Message 500, received by Customization Module 128, is a customization message (502-Yes), the message is checked to see that it contains or otherwise specifies a subscriber and appropriate subscriber profile data (504). If the Message 500 is a customization message, then it should contain essential subscriber identification information (subscriber ID) such as a password, to gain access to the system. If insufficient subscriber identification information is provided in the message (504-No), it is rejected and the procedure terminates. If the subscriber ID specified by the message is invalid, the message is rejected (506-No). If the subscriber ID is valid (506-Yes) and the subscriber profile data within message 500 passes any syntax or other verifications at step 502, the subscriber profile data is stored 508, preferably in a supplemental storage area such as a database dedicated to pending customer profile updates and/or registration requests. Alternatively, the pending subscriber profile data could be stored in an extra field of the subscriber record.

The subscriber is optionally notified of pending customization updates by any one of several ways (510). For example, a "subscriber customization validation" or "subscriber registration validation" message could be left in a designated voice mail box or E-mail address that corresponds to the subscriber.

Referring to FIG. 5B, when a subscriber logs on to the enhanced service platform (512), the customization module determines whether a customization update is pending (520). If there is a pending customization update (520-Yes), the Customization Module notifies the subscriber of the pending customization update and requests authorization from the subscriber to process the pending customization update (522). If there is no pending customization update, access to other subscriber functions (532) is initiated. The subscriber responds, either verbally or by keypad and that response is received by the Customization Module (524). If the subscriber grants authorization (526-Yes), the customization data corresponding to the pending customization update is registered in the subscriber's profile (528). If the subscriber does not grant permission (526-No), access to other subscriber functions (532) is initiated. A confirmation message is then sent to the subscriber's E-mail address or voice mail (530).

#### The Database Record

Referring to FIG. 6A there is shown an embodiment of a database record 600. Database record 600 represents the data structure of a valid Message 405 (FIG. 4A). Database record 600 may include data such as a subscriber identification 602, a transaction ID 604, a time stamp 606, which represents the time when transaction ID 604 was generated, and one or more parameters 608 to be stored in the subscriber's profile. Parameters 608 are subscriber related data such as an E-mail address indicating where incoming E-mail is to be forwarded, a pager telephone number, a fax number, a forwarding telephone number, etc.

Referring to FIG. 6B there is shown an embodiment of a database record 650. Database record 650 represents the data structure of a valid Message 500 (FIG. 5A). Database record

includes a subscriber identification 602 and one or more parameters 608 to be registered in the subscriber's profile.

#### Unified E-mail Address

Another preferred embodiment of the present invention is drawn to the case in which customization request 410 (FIG. 4A) or customization message 502 (FIG. 5A) is an E-mail message. In this preferred embodiment, the E-mail address that is used by a subscriber to send customization messages to the customization module is the same E-mail address that has been assigned to the subscriber for general E-mail services. Thus, the subscriber does not have to remember both his general E-mail address and a "customization" address. When the Enhanced Service Platform 100 receives an E-mail message that is addressed to a subscriber, the Platform automatically determines whether the E-mail message is a customization message. If the E-mail message is not a customization message, it is forwarded to the subscriber's inbox with no further processing. Customization E-mail messages are processed by the Customization module 128 according to methods such as those described in FIGS. 4 and 5. The determination of whether an E-mail message is a customization message, to be processed by customization module 128, or just general content for the subscriber can be effected by any number of methods, including a special designation in the "subject" line of the E-mail message or by a particular tag present in the E-mail message.

#### Alternate Embodiments

The present invention can be implemented as a computer program product that includes a computer program mechanism embedded in a computer readable storage medium. For instance, the computer program product could contain the program modules and managers shown in FIG. 1. These program modules and managers may be stored on a CD-ROM, magnetic disk storage product, or any other computer readable data or program storage product. The software modules and managers in the computer program product may also be distributed electronically, via the Internet or otherwise, by transmission of a computer data signal (in which the software modules are embedded) on a carrier wave.

While the present invention has been described with reference to a few specific embodiments, the description is illustrative of the invention and is not to be construed as limiting the invention. Various modifications may occur to those skilled in the art without departing from the true spirit and scope of the invention as defined by the appended claims.

What is claimed is:

1. A system for customizing a profile of a subscriber having access to a service platform comprising:

a subscriber service manager configured to receive a subscriber request at the service platform, the subscriber service manager upon receiving the subscriber request generating a transaction identifier, storing a copy of the transaction identifier, and communicating the transaction identifier to the subscriber; and

a customization module configured to receive and process a customization message sent to the system, the received customization message including subscriber profile data and a transaction identifier; the customization module, upon receipt of the customization message, comparing the transaction identifier in the received customization message with the copy of the transaction identifier stored by the subscriber service

## 11

manager, and registering the subscriber profile data when the transaction identifier and the copy of the transaction identifier correspond.

2. The system of claim 1, wherein the subscriber service manager is further configured to supply the subscriber with requested data from the profile when the subscriber makes a profile query request.

3. The system of claim 1, wherein the system is further configured to:

set a time-limit for receiving the customization message, the time-limit corresponding to a time that the transaction identifier was determined;

store in a database, in conjunction with the transaction identifier, the time-limit, and a subscriber identifier; and

reject the customization message if it is not received within the time-limit.

4. The system of claim 1, wherein the system is further configured to send a confirmation message to a destination corresponding to the subscriber when the subscriber profile data is registered.

5. The system of claim 4, wherein when the customization module receives a non-delivery notification indicating that the confirmation message was sent to an invalid or inoperative destination, the system is configured to use a telephone number corresponding to the subscriber to deliver to the subscriber a corresponding voice message.

6. The system of claim 1, wherein the subscriber profile data includes one or more subscriber-supplied parameters and when the customization module receives a customization message, the customization module is further configured to:

verify that each of the one or more subscriber-supplied parameters is within a corresponding range of values, and

notify the subscriber when a subscriber-supplied parameter is not within the corresponding range of values.

7. The system of claim 1, wherein the customization message that is sent to the system is routed through a mail server.

8. A method for customizing a profile of a subscriber having access to a service platform, the method comprising steps of:

receiving a subscriber request;

determining and storing a transaction identifier in response to the subscriber request;

communicating the transaction identifier to the subscriber;

receiving a customization message, the customization message including subscriber profile data and a transaction identifier;

comparing the stored copy of the transaction identifier with the transaction identifier in the received customization message; and

registering the subscriber profile data when the stored copy of the transaction identifier corresponds to the transaction identifier in the received customization message.

9. The method of claim 8, wherein when a profile query request from a subscriber is received, the method further comprises supplying the subscriber with requested data from the profile of the subscriber.

10. The method of claim 8, wherein the method further comprises steps of:

setting a time-limit for receiving the customization message, the time-limit

## 12

corresponding to a time that the transaction identifier was determined;

storing in a database, in conjunction with the transaction identifier, the time-limit, and a subscriber identifier; and

rejecting the customization message if it is not received within the time-limit.

11. The method of claim 8, wherein when the subscriber profile data is registered, the method further comprises sending a confirmation message to a destination corresponding to the subscriber.

12. The method of claim 11, further comprising the step of communicating a non-delivery voice message to the subscriber when a non-delivery notification is received in response to the step of sending the confirmation message to the destination corresponding to the subscriber.

13. The method of claim 8, wherein the subscriber profile data includes one or more subscriber-supplied parameters, the method further comprising the steps of:

verifying that each of the one or more subscriber-supplied parameters is within a corresponding range of values when a customization message is received, and

notifying the subscriber when a subscriber-supplied parameter is not within the corresponding range of values.

14. The method of claim 8, further comprising the step of routing the customization message through a mail server.

15. A computer program product for use in conjunction with a computer controlled service platform, the computer program product comprising a computer readable storage medium and a computer program mechanism embedded therein, the computer program mechanism comprising:

a subscriber service manager configured to receive subscriber requests at the service platform, the subscriber service manager upon receiving a subscriber request, determining a transaction identifier, storing a copy of the transaction identifier, and communicating the transaction identifier to the subscriber; and

a customization module configured to receive and process a customization message sent to the service platform, the received customization message including subscriber profile data and a transaction identifier; the customization module, upon receipt of the customization message, comparing the transaction identifier in the received customization message with the copy of the transaction identifier stored by the subscriber service manager, and registering the subscriber profile data when the transaction identifier and the copy of the transaction identifier in the received customization message correspond.

16. The computer program product of claim 15, wherein the subscriber service manager is further configured to provide the subscriber with data from the profile when the subscriber makes a profile query request.

17. The computer program product of claim 15, wherein the subscriber service manager is further configured to:

set a time-limit for receiving the customization message, the time-limit corresponding to the time that the transaction identifier was determined;

store in a database, in conjunction with the transaction identifier, the time-limit, and a subscriber identifier; and

reject the customization message if it is not received within the time-limit.

18. The computer program product of claim 15, wherein the customization module further configured to send a

## 13

confirmation message to a destination corresponding the subscriber when the customization data is registered.

19. The computer program product of claim 18, wherein when the customization module receives a non-delivery notification indicating that the confirmation message was sent to an invalid or inoperative destination, the customization module is configured to use a telephone number corresponding to the subscriber to deliver to the subscriber a corresponding voice message.

20. The computer program product of claim 15, wherein the subscriber profile data includes one or more subscriber-supplied parameters and when the customization module receives a customization message, the customization module is further configured to:

15 verify that each of the one or more subscriber-supplied parameters is within a corresponding range of values, and

notify the subscriber when a subscriber-supplied parameter is not within the corresponding range of values.

21. The computer program product of claim 15, wherein the customization message that is sent to the system is routed through a mail server.

22. A system for customizing a profile of a subscriber having access to a service platform comprising:

a customization module configured to receive a customization message sent to the system, the received customization message including subscriber profile data and a subscriber identifier that corresponds to a subscriber of the service platform; the customization module, upon receipt of the customization message, storing an indication of a pending subscriber customization request for the corresponding subscriber; and

35 a subscriber service manager configured to receive a request to access the service platform and, when the subscriber is accessing the system and there is a pending subscriber profile customization request for the subscriber, to request authorization from the subscriber to register the subscriber profile data specified by the previously received customization message, and to register the profile data upon receipt of such authorization.

23. The system of claim 22, wherein the subscriber service manager is further configured to supply the subscriber with requested data from the profile when the subscriber makes a profile query request.

24. The system of claim 22, wherein the subscriber profile data includes one or more subscriber-supplied parameters and when the customization module receives a customization message, the customization module is further configured to:

40 verify that each of the one or more subscriber-supplied parameters is within a corresponding range of values, and

55 notify the subscriber when a subscriber-supplied parameter is not within the corresponding range of values.

25. The system of claim 22, wherein the customization message that is sent to the system is routed through a mail server.

26. A method for customizing a profile of a subscriber having access to a service platform, the method comprising the steps of:

60 receiving a customization message, the customization message including subscriber profile data and a subscriber identifier that corresponds to a subscriber of the service platform;

## 14

storing an indication that the customization message has been received;

notifying the subscriber, when the subscriber is accessing the system, that the customization message has been received and requesting authorization from the subscriber to register the subscriber profile data; and registering the subscriber profile data upon receipt of such authorization.

27. The method of claim 26, wherein when a profile query request from a subscriber is received, the subscriber is supplied with requested data from the profile of the subscriber.

28. The method of claim 26, wherein the subscriber profile data includes one or more subscriber-supplied parameters, the method further comprising the steps of:

verifying that each of the one or more subscriber-supplied parameters is within a corresponding range of values when a customization message is received, and

notifying the subscriber when a subscriber-supplied parameter is not within the corresponding range of values.

29. The method of claim 26, further comprising the step of routing the customization message through a mail server.

30. A computer program product for use in conjunction with a computer controlled service platform, the computer program product comprising a computer readable storage medium and a computer program mechanism embedded therein, the computer program mechanism comprising:

a customization module configured to receive a customization message sent to the service platform, the received customization message including subscriber profile data and a subscriber identifier that identifies a subscriber of the service platform; the customization module, upon receipt of the customization message, storing an indication of a pending subscriber profile customization request for the subscriber; and

35 a subscriber service manager configured to receive a request to access the computer controlled service platform, and when the subscriber is accessing the service platform and there is a pending subscriber profile customization request for the subscriber, to request authorization from the subscriber to register the subscriber profile data specified by the previously received customization message, and to register the subscriber profile data upon receipt of such authorization.

31. The computer program product of claim 30, wherein the subscriber service manager is further configured to provide the subscriber with data from the profile when the subscriber makes a profile query request.

32. The computer program product of claim 30, wherein the subscriber profile data includes one or more subscriber-supplied parameters and when the customization module receives a customization message, the customization module is further configured to:

40 verify that each of the one or more subscriber-supplied parameters is within a corresponding range of values, and

notify the subscriber when a subscriber-supplied parameter is not within the corresponding range of values.

33. The computer program product of claim 30, wherein the customization message that is sent to the computer controlled service platform is routed through a mail server.

\* \* \* \*

3061

A standard linear barcode is positioned horizontally across the page, consisting of vertical black bars of varying widths on a white background.

US006269369B

**(12) United States Patent**

(10) Patent No.: US 6,269,369 B1  
(45) Date of Patent: Jul. 31, 2001

- (54) **NETWORKED PERSONAL CONTACT MANAGER**

(75) Inventor: **Brian D. Robertson**, Cambridge, MA (US)

(73) Assignee: **Amazon.Com Holdings, Inc.**, Seattle, WA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **08/962,997**

(22) Filed: **Nov. 2, 1997**

(51) Int. Cl.<sup>7</sup> ..... **G06F 3/00**

(52) U.S. Cl. ..... **707/10; 709/201; 709/202; 709/203; 709/300; 709/217; 709/218; 709/219; 707/1; 707/104; 707/501; 707/513; 707/201; 345/327; 345/331**

(58) Field of Search ..... **709/300, 201, 709/202, 203, 217, 218, 219; 707/1, 104, 501, 513, 10, 201; 345/327, 331**

(56) References Cited

U.S. PATENT DOCUMENTS

- |           |   |         |                       |        |
|-----------|---|---------|-----------------------|--------|
| 5,133,075 | * | 7/1992  | Risch .....           | 707/20 |
| 5,539,813 | * | 7/1996  | Jonsson et al. ....   | 379/20 |
| 5,737,726 |   | 4/1998  | Cameron et al. ....   | 705/   |
| 5,794,210 | * | 8/1998  | Goldhaber et al. .... | 717/51 |
| 5,796,395 | * | 8/1998  | De Hond .....         | 345/33 |
| 5,831,611 |   | 11/1998 | Kennedy et al. ....   | 345/33 |
| 5,835,087 | * | 11/1998 | Herz et al. ....      | 345/33 |
| 5,918,227 | * | 6/1999  | Polnerow et al. ....  | 707/10 |

#### OTHER PUBLICATIONS

- "1997 Edition GS Lotus Organizer/Exploring Organizer Windows 95 edition", pp. iii-vii and 8-1 – 8-6.  
"Applications Handbook for the Palm III™ Organizer" 3Com, pp. iii-v and 1-12, Apr. 1998.  
"Symantec ACT! User's Guide and Reference, Version 2.0 For Windows", pp. ix-xvi and 3-1 – 3-26.

"Information about Lotus Notes from the Lotus Web Site ([www.Lotus.com](http://www.Lotus.com))", pp. 1-13, Apr. 1998.

"Domino 4.6 Features Overview", Jul. 1997, 2 pages.  
[www.lotus.com/domino](http://www.lotus.com/domino).

"Yahoo! White Pages", Copyright 1994-1998, Yahoo!,  
<http://www.yahoo.com/>.

"WhoWhere?", Copyright 1995-1998, WhoWhere? Inc.  
<http://www.whowhere.com/>.

Randall A. Tamura et al "Lotus Notes 4" pp. 39, 64-65, and 176, 1996.\*

\* cited by examinee

*Primary Examiner*—Dung C. Dinh

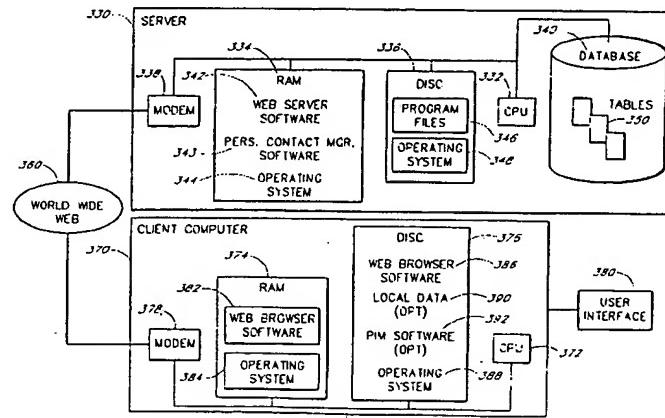
*Assistant Examiner—Hieu C. Le*

(74) *Attorney, Agent, or Firm*—Knobbe, Martens, Olson & Bear, LLP.

(57) ABSTRACT

A network-computer-based personal contact manager system is disclosed wherein users of networked clients maintain and update a set of user information which is stored in a relational database on a networked server. The personal contact manager system allows each user to specify on an individual basis which of their contacts are permitted to access respective datums of their user information. In some cases, and assuming permission is granted, the system will issue notifications (e.g., by e-mail) to a user's contacts when the user changes his information or when a preset event, such as a birthday, as defined by the user, is to occur. The system also allows users to find contacts based on common group affiliations and notifies users when there are coincidences in their data (e.g., travel plans, astrological compatibility). The personal contact manager system supports the retrieval of information on the contacts of contacts, assuming such as permission has been granted by the contacts and their contacts, and can also be used to synchronize the server database with a PIM database of the user and any contacts of the user who have the appropriate permissions.

16 Claims, 14 Drawing Sheets



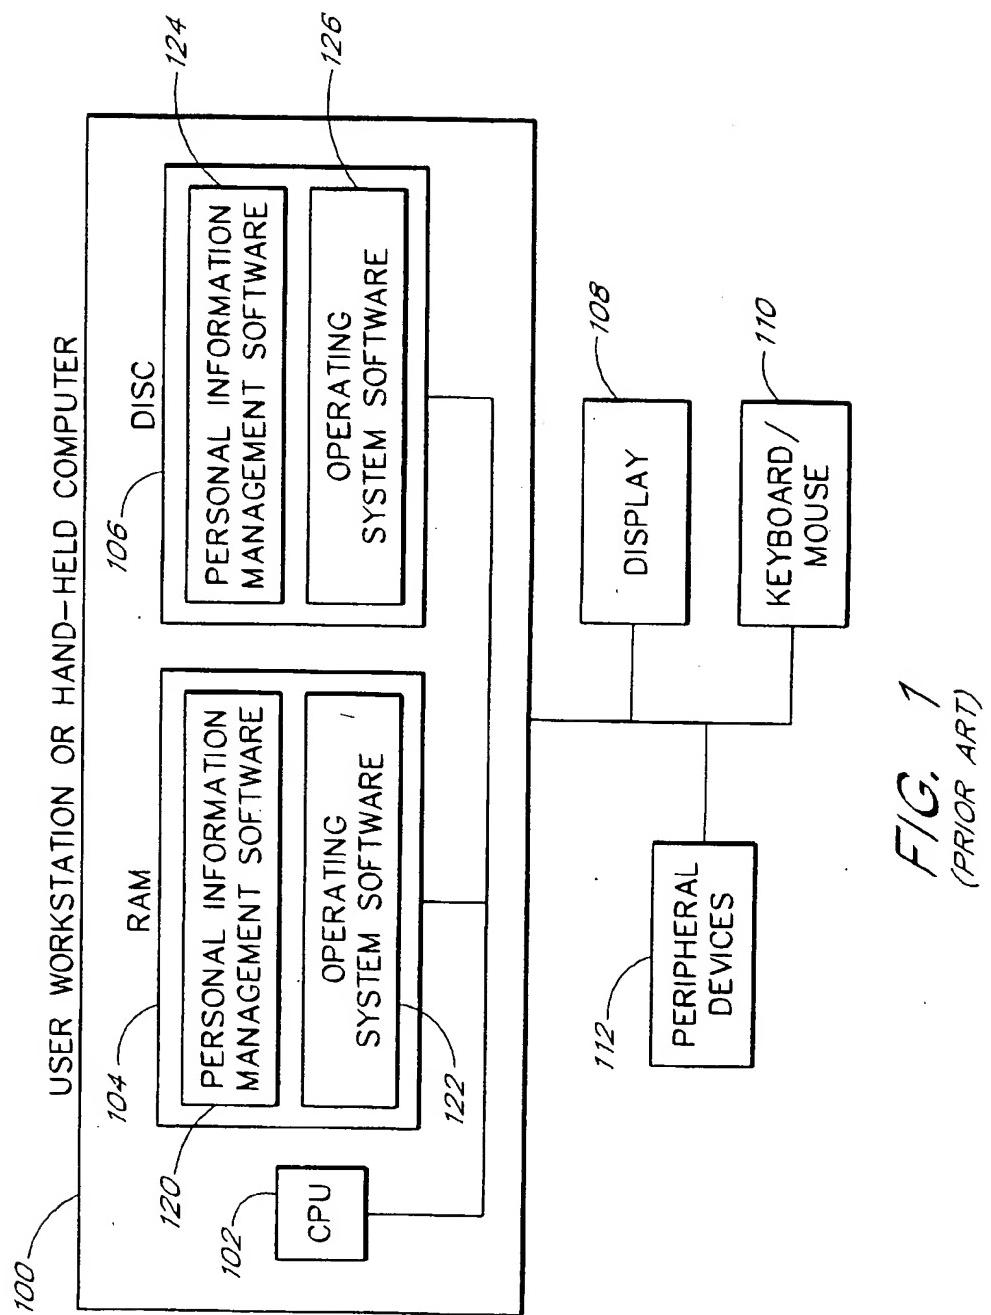


FIG. 1  
(PRIOR ART)

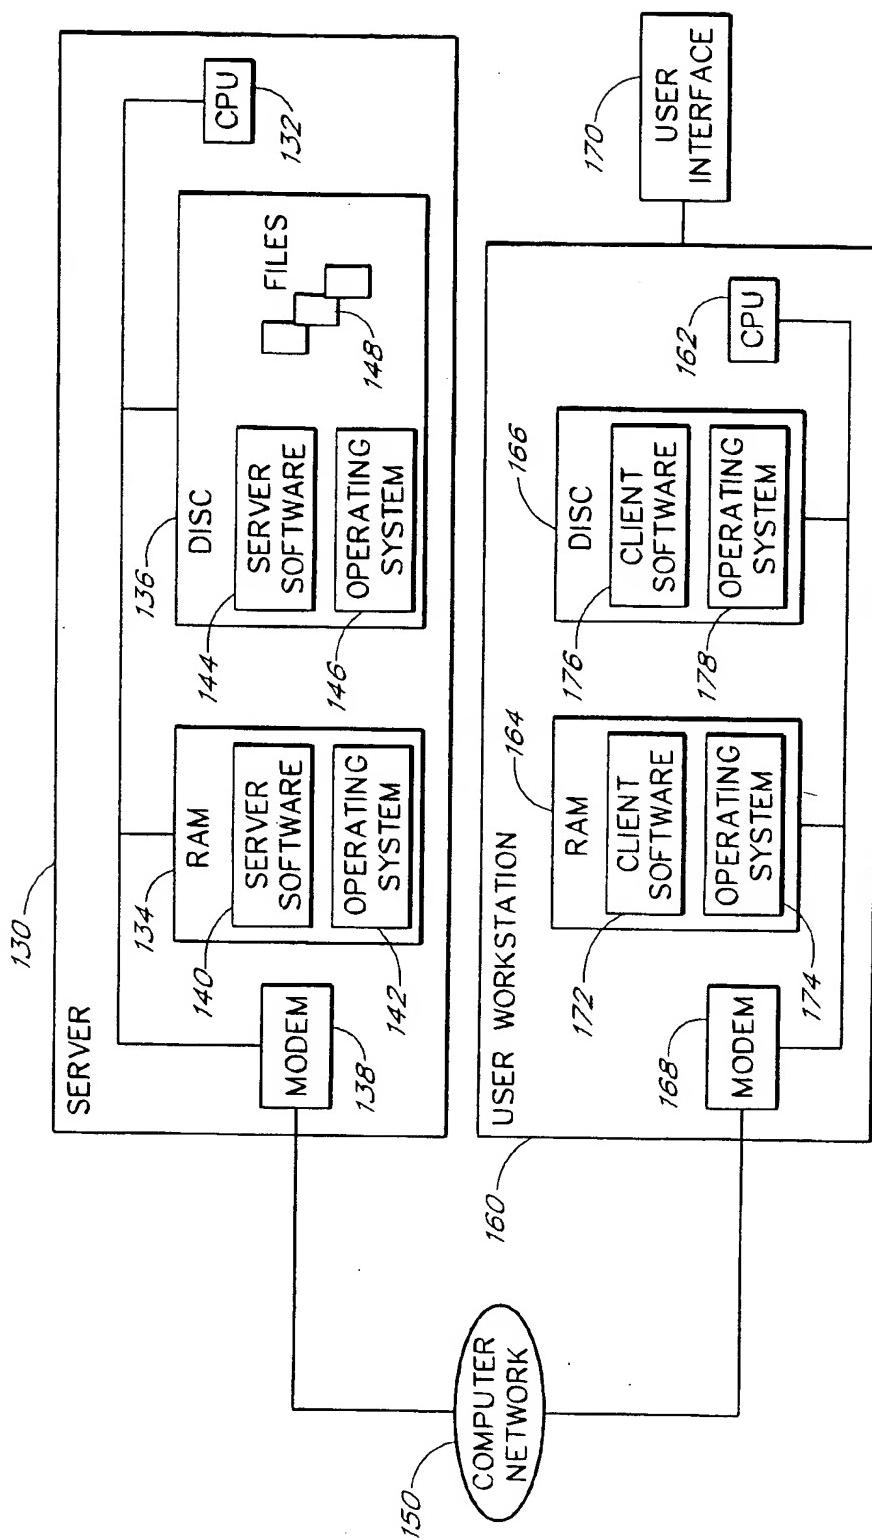


FIG. 2  
(PRIOR ART)

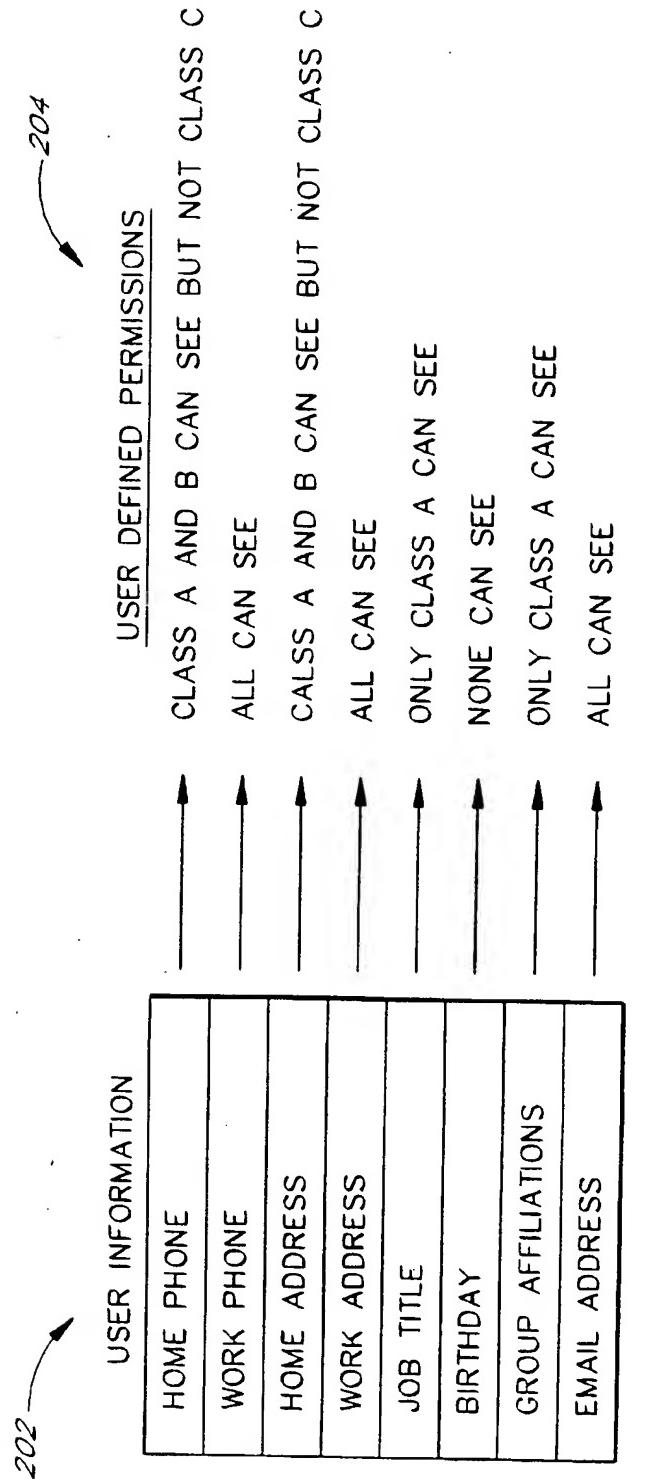
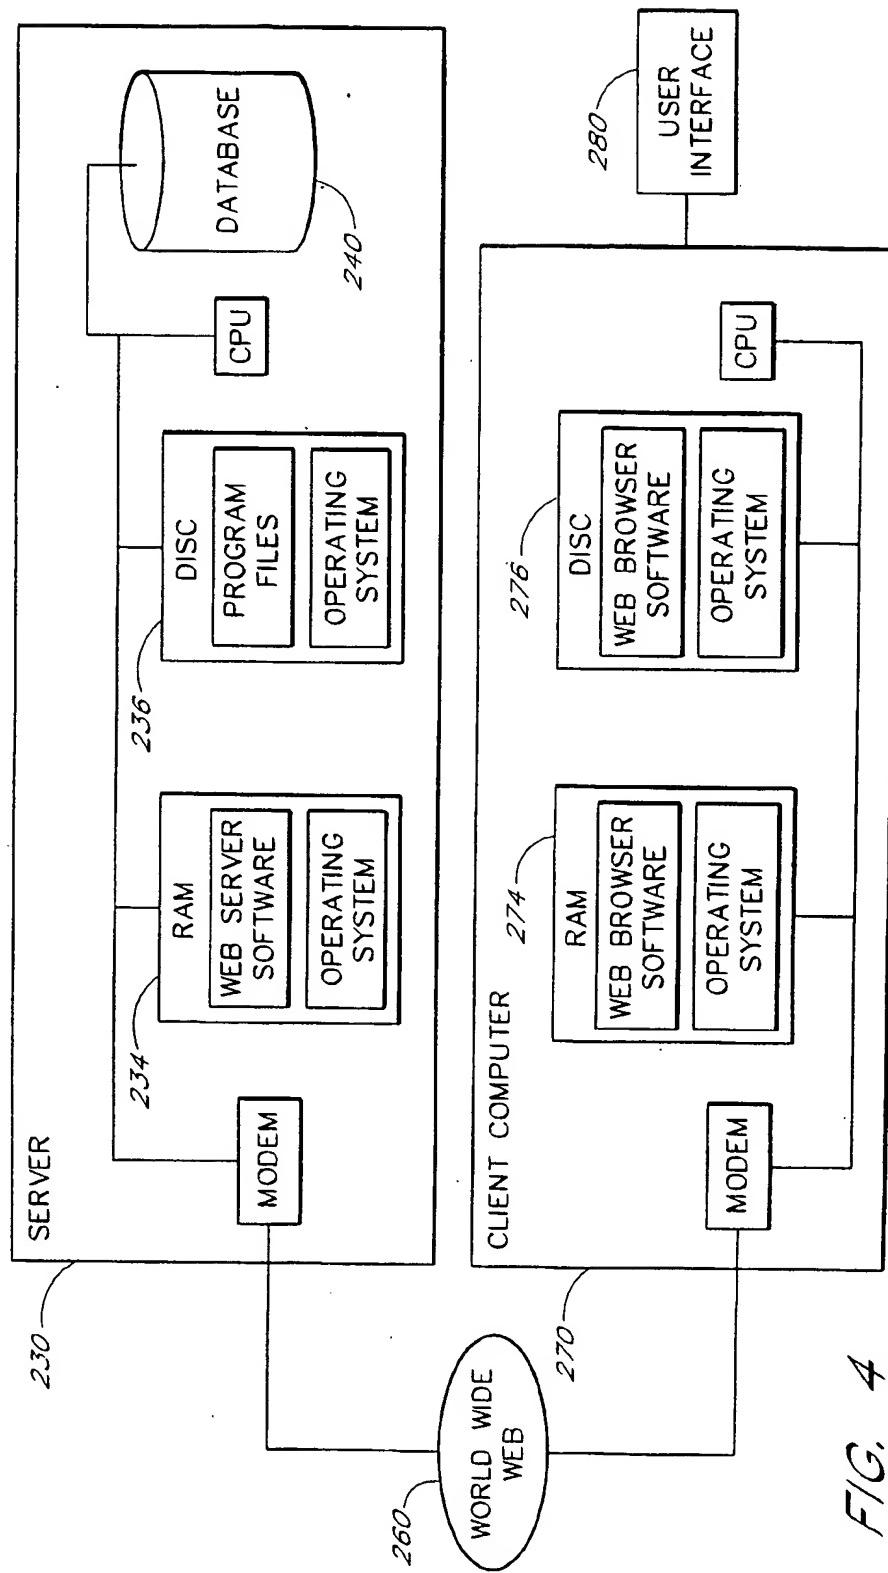


FIG. 3  
(PRIOR ART)



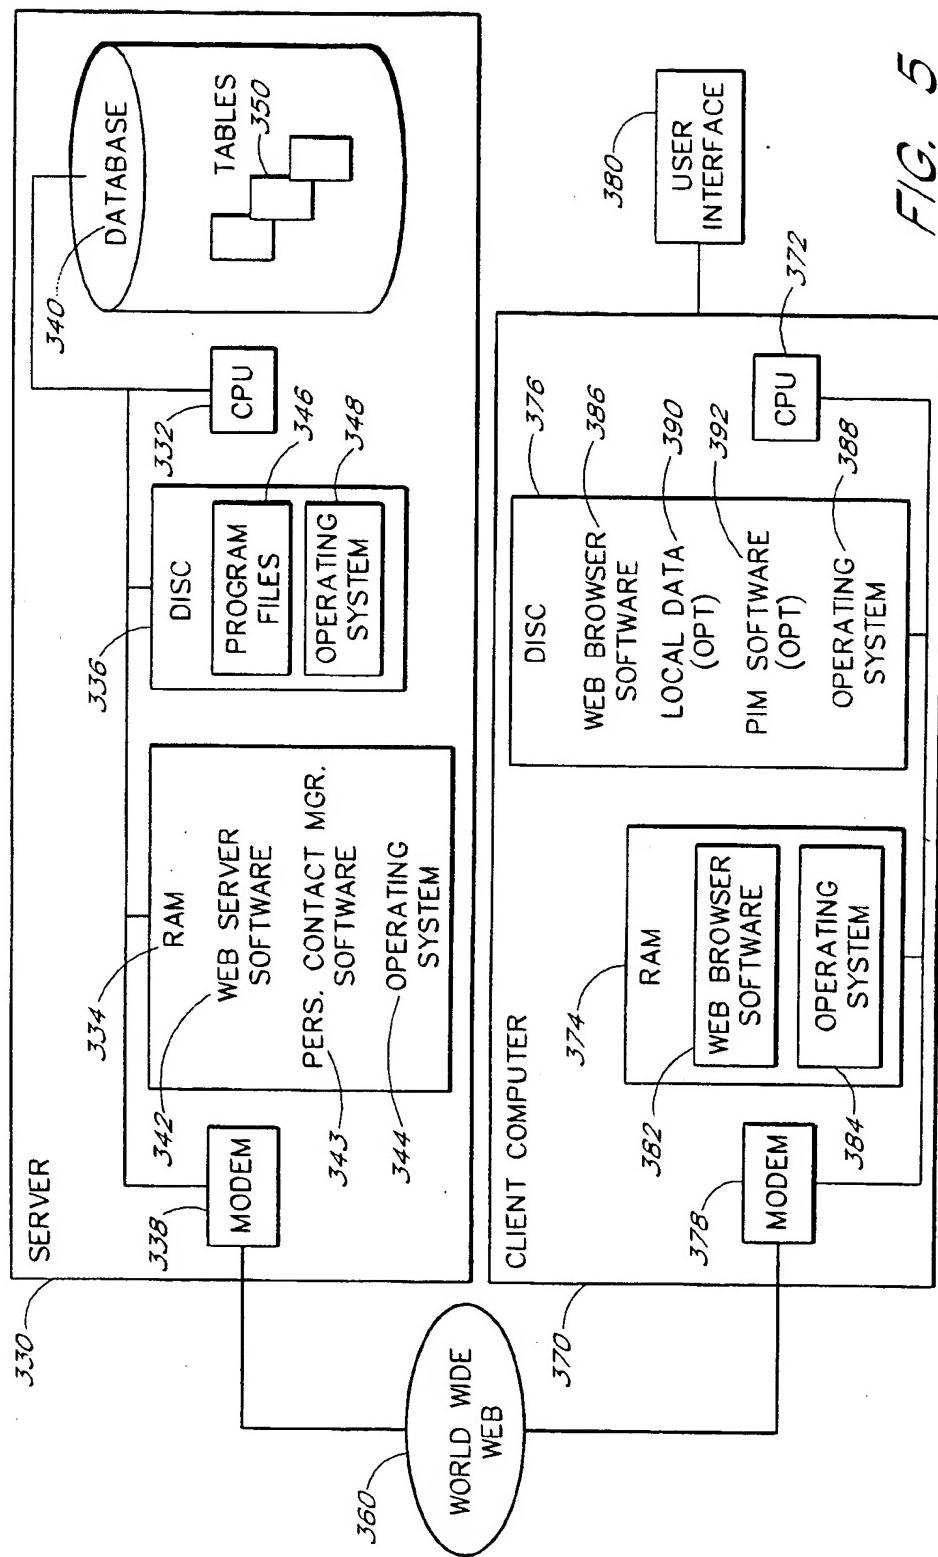


FIG. 5

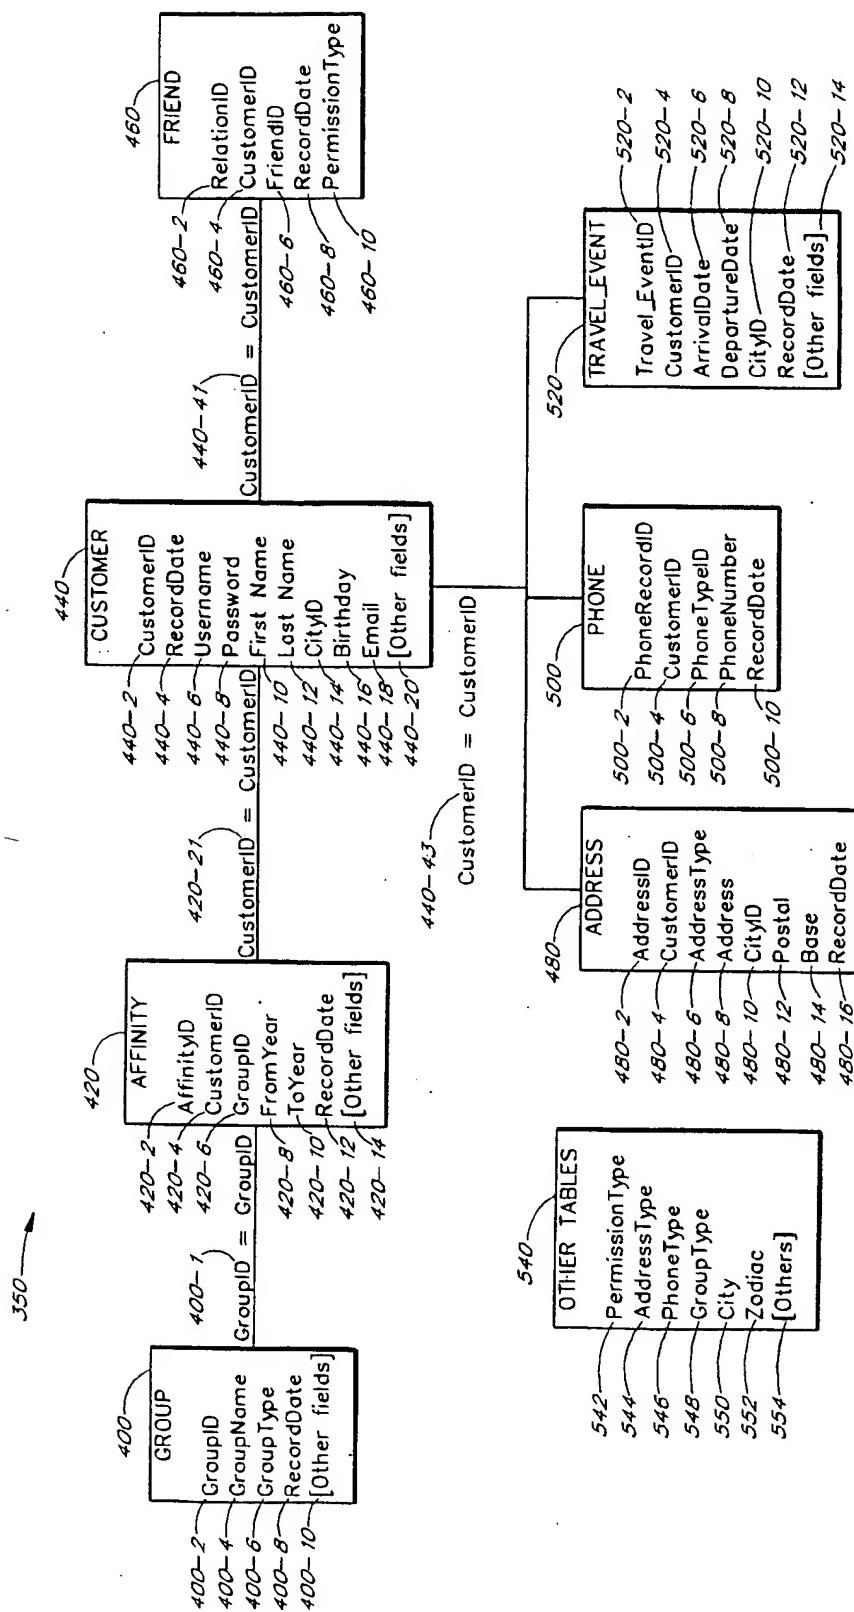


FIG. 6

560 →

## PSEUDO REGISTRATION FORM

Name	560-2
Home Address	560-4
Home Phone	560-6
Work Address	560-8
Work Phone	560-10
Birthday	560-12
Your High School	560-14
Year of Enrollment	560-16
Graduation Year	560-18
Your College	560-20
Year of Enrollment	560-22
Graduation Year	560-24
Submit	560-26

FIG. 7

580 →

## PSEUDO GROUP LIST FORM

580-2 →

Following are the other members who went to  
your college at about the same time.

580-4 →

Click on the boxes next to the names of the people  
you'd like to add to your Address Book.

580-6 →

State University, 1982-1986

580-8 →

580-14 →



John Doe (Graduated 1985)

580-12 →



Robert Johnson (Graduated 1986)



Jane Smith (Graduated 1986)

580-16 →

Submit

FIG. 8

600 →

## PSEUDO PERMISSION FORM

600-2 →

Click on the boxes next to the Permission Levels that you would like to grant to your new contact

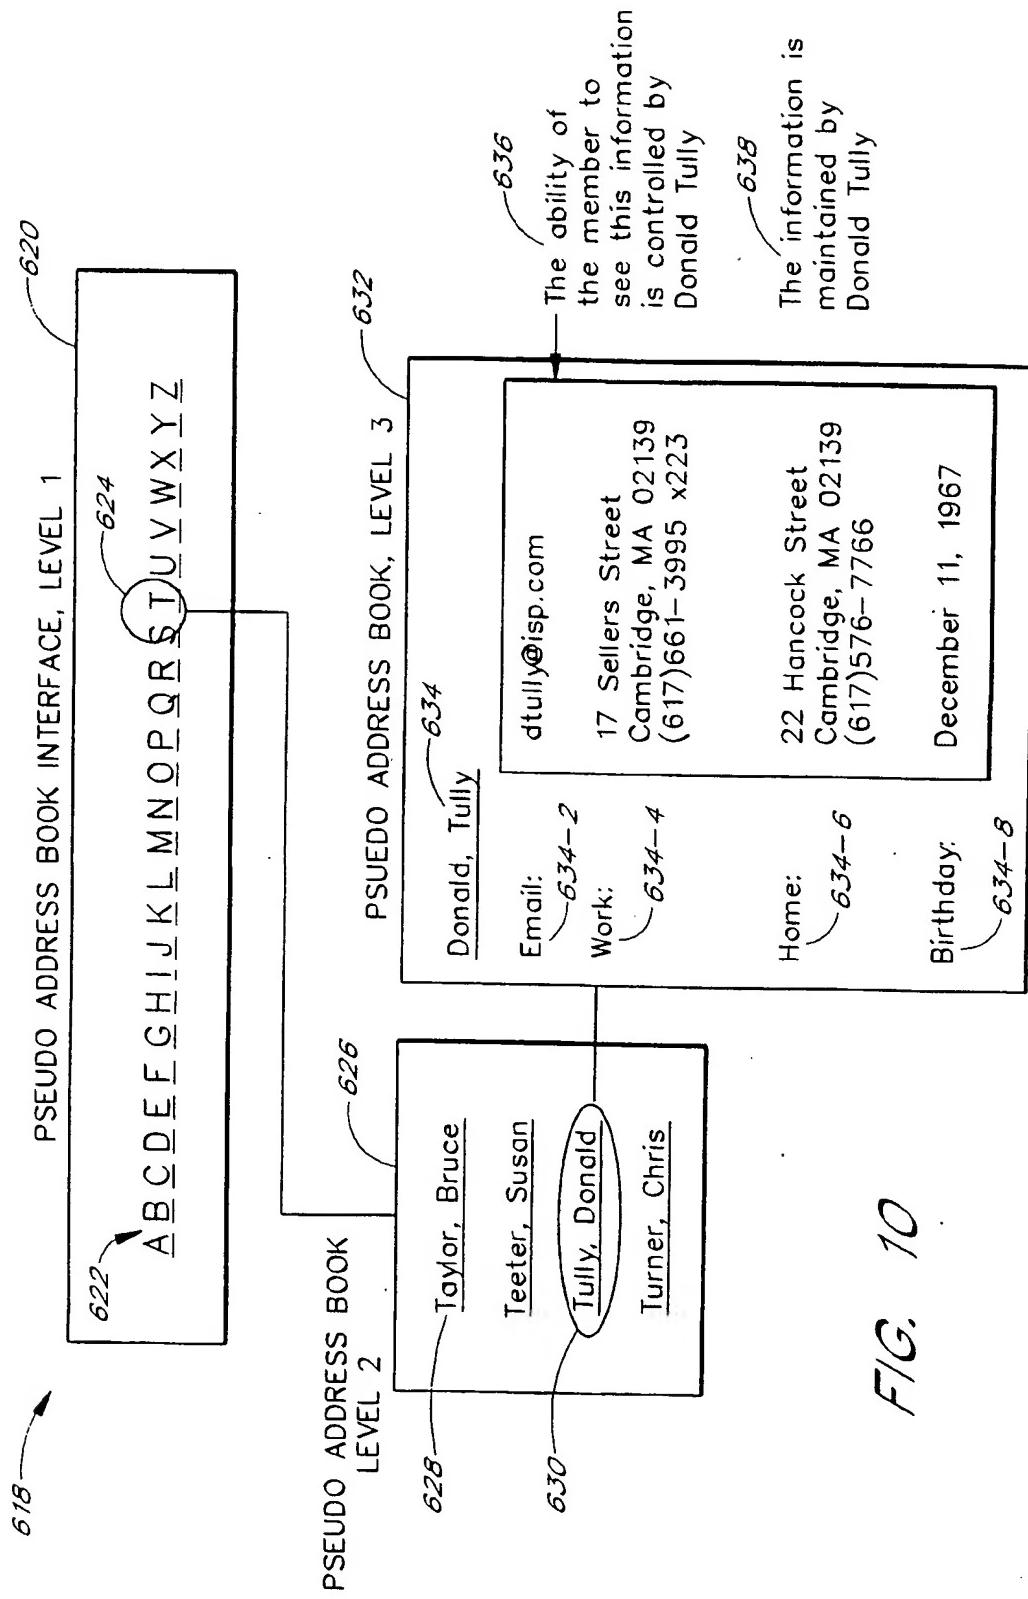
600-4 →  
John Doe

- 600-7 →  Crossing Paths Notification Permission  
600-9 →  Personal Information  
600-10 →  Work Information  
600-12 →  Birthday Notification  
600-14 →  Friend of Friends Information

600-16 →

Submit

FIG. 9



F/G. 10

650 —

## PSEUDO MEMBER UPDATE

Member Update 650-2  
December 7, 1998

The following of your contacts have upcoming birthdays:  
— Avery Rogers (Dec. 11) } 650-6  
— Jane Bigelow (Dec. 14) } 650-4

Your contacts have registered the following address changes:

— Tom Kohn

New work address:  
1000 Wilson Boulevard } 650-10  
Arlington, Va 22229 }

New work phone: } 650-12  
703-558-3312 }

The following new members have affiliated with the same groups as you:

— Gary Clayton (State College, 1985)

650-14

The following members have linked to you and have requested that  
you reciprocate:

— Jun Obama }  
— Lee Rogers } 650-18

You have scheduled a trip to Phoenix on December 14. The following of  
your contacts live in Phoenix or will be in Phoenix on that date:

— Andrew Kress } 650-22  
— Taylor Pierce }

According to Astrology, the following of your contacts are compatible with  
you today (Libras)

— Bryan Jamieson } 650-26  
— Anne Thierry }

FIG. 11

PSEUDO ADD TRAVEL FORM

660-1  
Traveling To City: [ ]

660-2  
State: [ ]

660-4  
Country: [ ]

660-6  
Arrive in City on Date: [ ]

660-8  
Leave City on Date: [ ]

660-10  
How to get in touch  
while in this City: [ ]

PSEUDO CROSSING PATH LIST

670-2  
The following of your contacts will be in Chicago  
while you are there.

670-4  
Check the boxes next to their names if you would like  
to inform them that you will be visiting Chicago.

670-6  
People Who Live in Chicago

Scott Ullerm

Taylor Pierce

Betsy Klein

670-8  
People Who Will be Visiting Chicago

Tania Gutsche

670-10  
Submit

F/G. 12

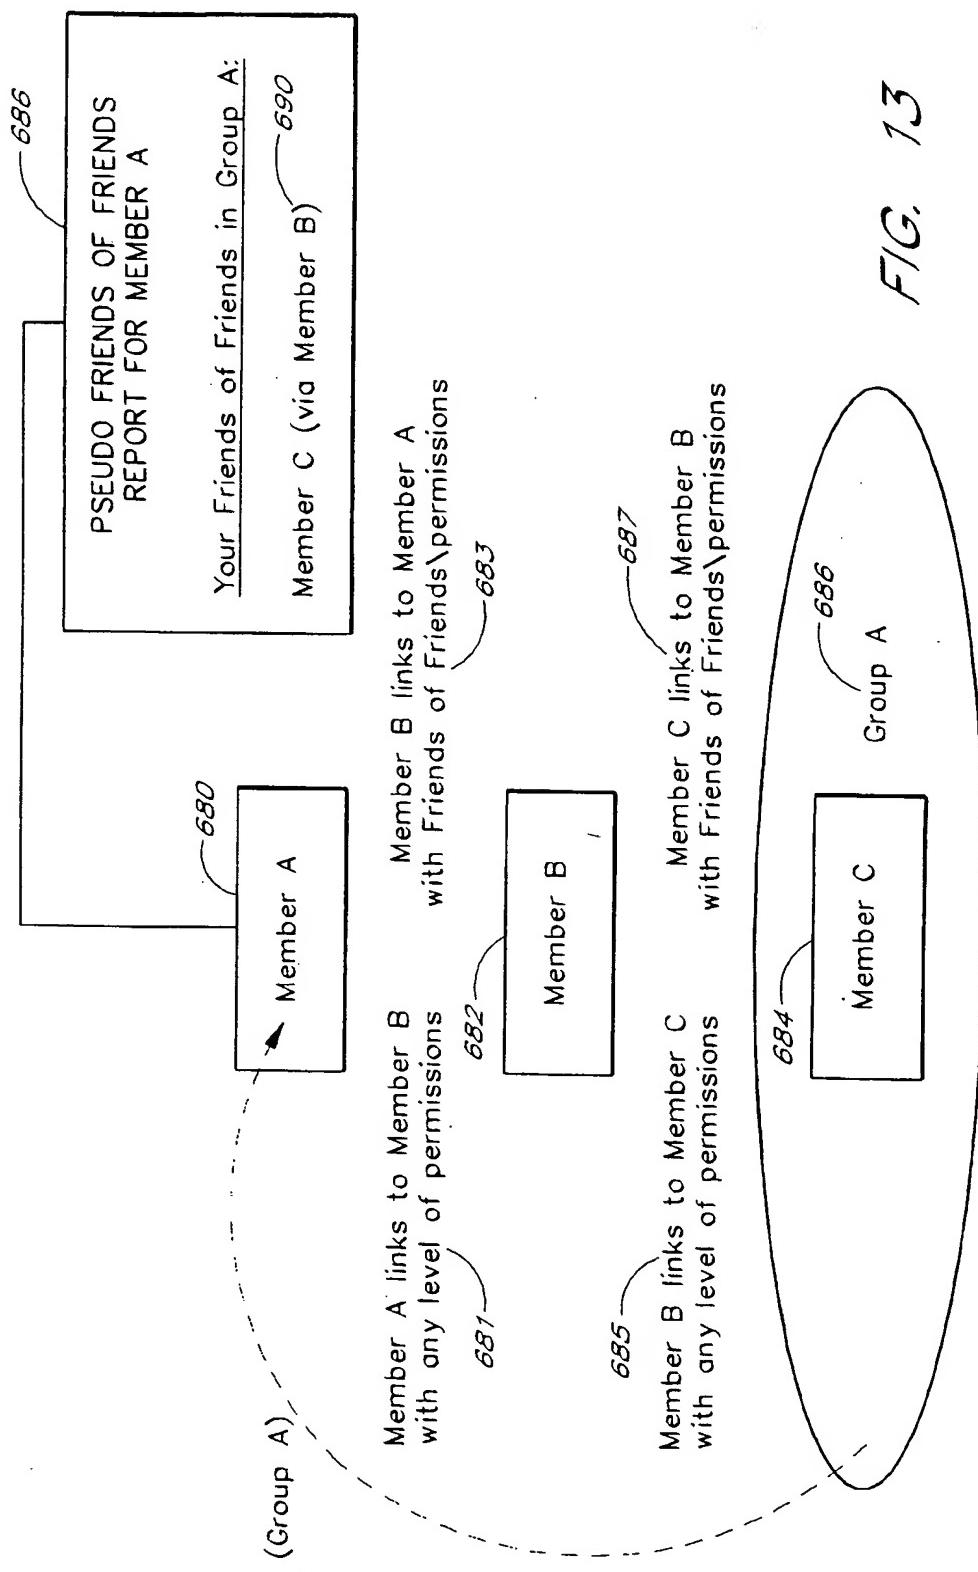


FIG. 13

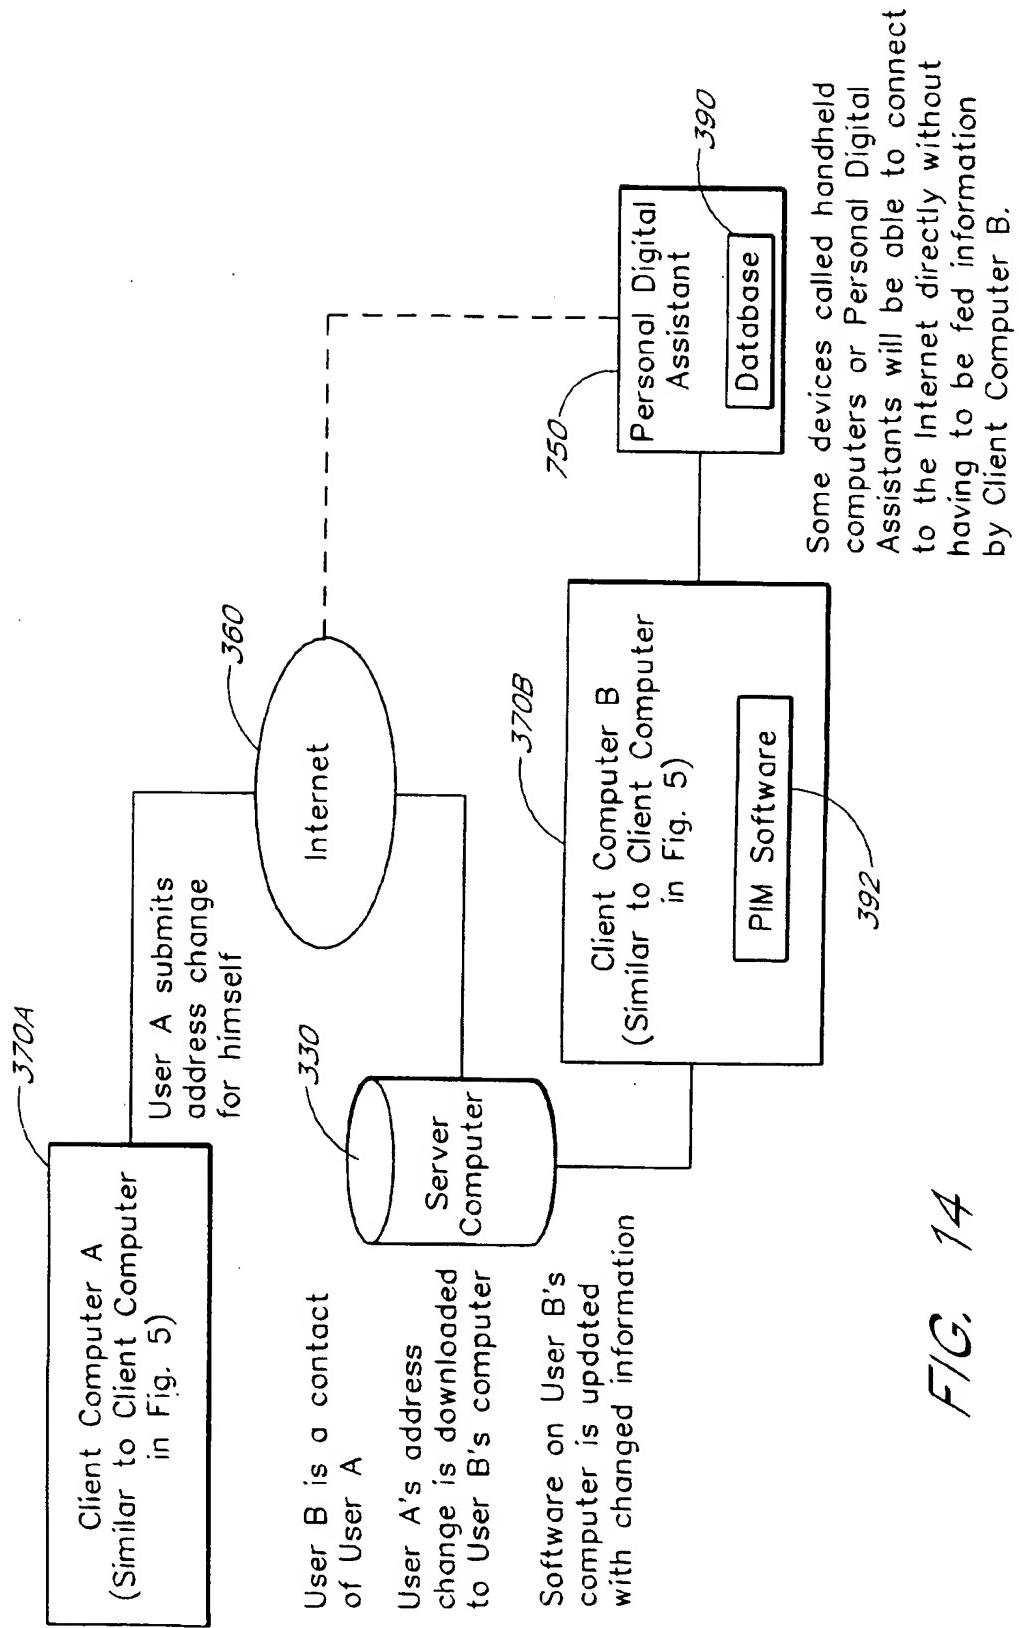


FIG. 14

## NETWORKED PERSONAL CONTACT MANAGER

The present invention relates generally to computer software used to manage contact information—such as mailing addresses, e-mail addresses, phone numbers, and birthdays—and more specifically to a method of creating links between members over a network and providing information to each member based on levels of permission maintained by the other members to which they are linked.

### BACKGROUND OF THE INVENTION

Several types of prior art for managing contact information exist, including Personal Information Management software applications, Groupware Applications, and Internet-based "White Pages" and e-mail services.

#### Personal Information Management Software

As represented generally in FIG. 1, in a typical prior art Personal Information Management (PIM) software application (e.g., Lotus Organizer, Microsoft Outlook, or U.S. Robotics Palm Pilot), a PIM software application 120, 124 that stores contact information in a database resides on a workstation or handheld computer 100 having a central processing unit 102, a display 108, a keyboard and/or mouse 110, a primary memory 104 (e.g., random access memory) for program execution, a secondary memory 106 (e.g., a hard disc) for program storage, and peripheral devices 112. As is well known, programs, such as the PIM software 120, are executed in the RAM 104 by the CPU 102 under control of the operating system software 122, 126.

In the prior art, users themselves enter the contact information that they want to store in the PIM software. A variety of methods exist for entering this contact information. It may be entered manually using the keyboard, imported from an existing file on their computer, or imported via a peripheral device such as a business card scanner. The defining characteristic of this class of prior art is that the input of the contact information is performed by the user of the software and, when the information changes, the user must modify the information himself. What this class of prior art lacks is a means for information to be shared between multiple users and a means for a given user to post changes to his own information for the benefit of others.

#### Groupware Applications

As generally represented in FIG. 2, in a typical prior art Groupware application (e.g., Lotus Notes), a user workstation 160 accesses information stored on a central server computer 130 over a computer network 150, such as a Local Area Network or Intranet. The server system consists of a central processing unit 132, a primary memory 134 (e.g., random access memory) for program execution, a secondary storage device 136 (e.g., a hard disc) for program storage, and a modem 138 or other device for connecting to the computer network. The user workstation 160 is the same as the user workstation 100 described in reference to FIG. 1 with the addition of a modem 162 or other device for connecting to the computer network. The file server or database contains data files 148 that can be accessed only by authorized users. The user uses client software 174, 176 running on the user workstation 160 to access the files 148 under the mediation of server software 140, 144 running on the server 130.

Typically, in such a system a central system administrator organizes users into classes and the creator of a file 148 determines what classes of users may view the file. The rules governing which individual users or classes of users have the authorization to view a particular file 148 may be stored

as part of the file itself. Alternatively, these rules are based upon the hierarchical directory structure of the file server in which the file is stored. That is, a particular user may view files in one directory but not another.

FIG. 3 represents a common deployment of a contact management system based on Groupware. Each user enters information 202 about himself and specifies a set of permissions 204 that define what classes of users are able to view various pieces of the information 202. What this deployment of the prior art lacks is the ability to authorize viewing privileges on a user-by-user basis rather than on a class-by-class basis. For instance, a user would be able to grant access to his home phone number 206 to the Human Resources department of his employer (e.g., Class A) while denying access to the same information to his co-workers (e.g., Class C). The user would not be able to give access to his home phone number selectively to a first co-worker while denying it to a second co-worker if both co-workers were part of the same class of users as organized by the central system administrator. Furthermore, such a system would lack a practical notification methodology. There would be no way for a user to specify "notify me when the first co-worker changes his information but not when the second co-worker changes his information."

#### Internet-Based "White Pages" and E-Mail Directory Services

In a typical prior art "white pages" or e-mail service, client computers and a server computer are connected via the World Wide Web as depicted in FIG. 4. A user subscribes to a White Pages or E-Mail service via a client computer 270 operating a web browser 282 or other software application residing in memory 274 that allows it to display information downloaded from a server computer 230 over the World Wide Web 260. The server computer system accesses a database 240 containing contact information entered by registered users. The service enables users to view contact information entered by other users. The authorization scheme may allow all users to limit certain classes of users from viewing certain parts of their user record as represented in FIG. 3. However, there are no linkages between individual users and thus users cannot restrict the viewing of their information on a user-by-user basis. Furthermore, users cannot be notified when information for particular users has changed.

### SUMMARY OF THE INVENTION

The present invention is a computer-network-based contact management system that allows members to create and maintain contact with other members and determine on a person-by-person basis what information to share or withhold. The system is based on a relational database scalable to millions of users that resides on a server computer.

The invention was developed shortly after the advent of the World Wide Web, which promoted millions of people worldwide to connect their computers with a standard protocol, a phenomenon which made the invention practical and beneficial.

When a user becomes a member of the system, the member associates himself with any number of affinity groups and creates a data record for himself by entering information in specific data fields. Based on the affinity groups with which the user has associated himself, the system then informs the user of other members in the same groups and allows the user to establish a link to any of those members on an individual basis.

For each second user to which a first user has established a link, the second user can specify which data fields in his

data record can be viewed by the first user. Each second user to which the first user has established a link is informed that a link has been established. The second user can in turn decide whether or not to establish a link to the first user. If the second user chooses to establish a link to the first user, he can specify which data fields in his personal data record can be viewed by the first user. In addition, each time a new user associates himself with a group with which an existing user has associated himself, the present invention informs the existing user that the new user has joined that group and allows the existing member to establish a link to the new user.

For each first user, the present invention maintains a database of information about the second users to whom the first user has established a link. The personal address book of the first user contains the information in the data fields that the second users have given the first user permission to view. Whenever a second user changes any information in any data field of his data record, the information in that field is automatically updated in the information database of each first user whom he has given permission to view the information in that data field.

In addition, if a first user has given a second user the proper form of data field permission for the personal data record of the first user, the present invention will inform the second user whenever first user's birthday or anniversary is approaching, whenever the first user will be travelling in the vicinity of the second user, and whenever the astrological sign of the first user is compatible with the astrological sign of the second user.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and form a part of this specification, illustrate embodiments of the invention and, together with the description, serve to explain the principles of the invention, wherein:

FIG. 1 depicts a computer loaded with Personal Information Management software;

FIG. 2 generally depicts the data schema of a category of prior art known as groupware applications;

FIG. 3 shows a common scheme for authorizing permission to view information in the prior art;

FIG. 4 depicts two computers interconnected via the Internet, one of which is a server connected to a database and the other of which represents a user's client workstation, both of which are configured according to the prior art;

FIG. 5 depicts two computers interconnected via the Internet, one of which is a server connected to a database and the other of which represents a user's client workstation, both of which are configured according to the present invention;

FIG. 6 represents an object model of the key tables in the relational database maintained on the server computer in the preferred embodiment of the present invention;

FIG. 7 represents a pseudo graphical user interface in which a user enters information in specific data fields to create a personal data record;

FIG. 8 represents a pseudo graphical user interface for listing other users with the same group affiliation as that specified by a first user;

FIG. 9 represents a pseudo graphical user interface for specifying what type of data fields from a first user's personal data record to which the first user wishes to grant a specific second user access;

FIG. 10 represents a pseudo graphical user interface that displays the information stored in a user's personal address book;

FIG. 11 represents a pseudo graphical user interface that provides a first user with specific information that has changed about the other users to which the first user is linked;

FIG. 12 represents a pseudo graphical user interface that allows a first user to enter travel information and find out which contacts have overlapping travel schedules

FIG. 13 represents a pseudo graphical user interface that allows a first user to gather information about the contacts of his contacts; and

FIG. 14 is a data flow diagram of an alternative embodiment of the present invention where a personal digital assistant is synchronized with a server database of user information.

#### DESCRIPTION OF THE PREFERRED EMBODIMENT

Reference will now be made in detail to the preferred embodiments of the invention, examples of which are illustrated in the accompanying figures. While the invention will be described in conjunction with the preferred embodiments, it will be understood that they are not intended to limit the invention to those embodiments. On the contrary, the invention is intended to cover alternatives, modifications and equivalents, which may be included within the spirit and scope of the invention as defined by the appended claims.

As represented in FIG. 5, the preferred embodiment follows a standard Internet architecture, in which client computers 370 and a server computer 330 are connected via the World Wide Web 360 and modems 338, 378 or other communications channels. A user accesses the server 360 via a client computer 370 operating a web browser 382 or other software application residing in memory 374 that allows it to display information downloaded from a server computer 330. The server computer system 330 runs server software 342, including the network-computer-based personal contact manager 343 of the present invention, which interacts with the client computers 370 and a user information database 340. In a commercial embodiment of the present invention, the personal contact manager 343 is the heart of a Web-based personal contact management service called PlanetAll. The database 340 contains contact information entered by registered users. The personal contact manager 343 in some situations will notify a set of users of updates made to the database 340 by another user to whom the notified set is related.

A distinction over the prior art is that the database 340 in the present invention is necessarily a relational database built from a set of relational tables 350. In the conventional manner, both the server 330 and the clients 370 include respective storage devices, such as hard disks 336 and 376 and operate under the control of operating systems 344, 384 executed in RAM 334, 374 by the CPUs 332, 372. The server storage device 336 stores program files 346 and the operating system 348. Similarly, the client storage devices 376 store the web browser software 386 and the operating systems 388. In an alternative configuration, in which the client is a personal information manager (PIM), such as the U.S. Robotics Palm Pilot, the disc 376 can also include a local PIM database 390 and PIM software, which performs data management and synchronization functions.

FIG. 6 outlines the data structure of the relational database 340 in the preferred embodiment, in which seven tables 350 are employed to enable most of the functionality of the present invention:

- (1) Customer Table 440;
- (2) Friend Table 460;
- (3) Group Table 400;

- (4) Affinity Table 420;
- (5) Address Table 480;
- (6) Phone Table 500; and
- (7) Travel Event Table 520;

The Customer Table 440 contains one record for each unique user. The key field in this table is CustomerID 440-2. All information stored in the various database tables relating to a particular member is linked together by a unique number in this field. Other important fields in this table include information used by users to login to the system (Username 440-6 and Password 440-8), information which helps users identify each other (First Name 440-10, Last Name 440-12, and E-mail 440-20), information required to provide Birthday Notification (Birthday 440-16) and information required to provide Crossing Paths notification (CityID 440-14). Each record in the Customer Table 440 is time-stamped via the RecordDate field 440-4. Other fields 440-22 can also be included in the Customer Table 440 (and the other tables as well).

The Friend Table 460 is a key to the present invention because it relates users to each other. Each record in the table represents a relationship between one user, identified by CustomerID 460-4, and another, identified by FriendID 460-6, with a certain level of permissions 460-10. The user interface of the present invention provides a multitude of ways for users to view information about other users, and every one of these ways relies on a database query of the Friend Table 460 to determine the list of other users whose information a particular user may see. Each record is time-stamped via the RecordDate field 460-8 so that users may be notified when their contacts' records change. Each record is uniquely identified by a RelationID 460-2.

The Group Table 400 contains one record for each unique group with which users may affiliate. Each group is identified by a GroupName 400-4 and GroupType 400-6. Examples of these groups would be GroupName 400-4—"Massachusetts Institute of Technology" (GroupType—"University") and GroupName 400-4—"Sigma Chi" (GroupType—"Fraternity"). Each record has a time-stamp 400-8 and a unique identifier 400-2.

Each record of the Affinity Table 420 relates a user, identified by CustomerID 420-4, to a group, identified by GroupID 420-6. If a user affiliates with six groups, there would be six records in the Affinity Table 420. This table stores information about the time period of a user's affiliation with a particular group in the FromYear and ToYear fields 420-8, 420-10 so that the system may help users find their contemporaries. Each record is time-stamped 420-12 so that the system may report to users when other users join the group, has a unique identifier 420-2 and can include additional fields 420-14.

The Address Table 480 stores information for any number and kind of addresses for a particular user, identified by CustomerID 480-4. For instance, if a user wants to make his home address, work address and summer home address available to his contacts, there would be three records for that user in the Address Table 480, each being identified in part by an appropriate AddressType 480-8 (e.g., home, work, summer home). Each record is time-stamped 480-16 so that the system can notify users when their contacts have added or modified address information and has a unique identifier 480-2. Address information is conventional, including street Address 480-8, CityID 480-10, Postal code 480-12, and military Base 480-14 fields.

The Phone Table 500 is directly analogous to the Address Table 480, but it stores telephone and fax number information instead of address information. Each record is identified

by a unique PhoneRecordID 500-2 and includes the CustomerID 500-4 of the user whose phone information is contained in the record, a phone type ID 500-6 indicating, e.g., whether the record is for a telephone or fax, the phone number 500-8 and a time-stamp 500-10.

The Travel Event Table 520 stores information about users' travel plans. This table is required to notify users when their travel plans intersect with the travel plans of their contacts. A record in the Travel Event Table 520 includes the CustomerID 520-4 of the user whose travel information is contained in the record, arrival and departure dates 520-6, 520-8 and a CityID 520-10 identifying the travel destination. Each record is uniquely identified by a Travel\_EventID 520-2 and is time-stamped with a RecordDate 520-14.

In the preferred embodiment, a multitude of other tables 540 are used to enable a variety of user services. The Permission Type Table 542 contains one record for each of the varieties of permission levels the system allows members to assign to their contacts in the Friend Table 460. In the preferred embodiment, as illustrated in FIG. 9, permission information is grouped into five categories for the purpose of user interface simplicity (crossing paths notification permission 600-6, personal information 600-8, work information 600-10, birthday notification 600-12, and friends of friends information 600-14). However, the Permission Type table 542 could just as easily be structured to allow members to grant and deny access to information on a field by field basis.

The City Table 550 stores latitude and longitude information for two million cities to enable the system to notify users when their contacts travel within a defined geographical radius. The Zodiac Table 552 allows the system to associate birthdays with signs of the Zodiac and thereby notify which of their contacts have compatible astrological signs on a particular day. The AddressType, PhoneType and GroupType tables 544, 546, 548 define the types of address, phone and group that can be defined in the respective Address, Group and Phone tables 480, 400, 500. The advantage of this normalized relational database architecture is that it permits scaling and speed far in excess of any embodiment of the prior art.

FIGS. 7 through 12 display pseudo software graphical user interfaces (GUIs). In the preferred embodiment, the web server software 342 on the server computer 330 displays these GUIs via the computer communications interface 360 on the user interface 380 of the user workstation computer 370. The database and communications operations necessary to perform the described functions are controlled by the personal contact manager 343, which employs where necessary the services of the web server software 342. For example, the personal contact manager 343 updates the database tables 350 when a user submits a new home address and then determines whether any of that user's contacts need to be notified of the change. If so, the personal contact manager 343 will issue the notifications via the web server software 342. It should be assumed, unless a statement to the contrary is made, that all of the operations described herein which are aspects of the present invention are embodied by the personal contact manager 343.

Referring now to FIG. 7, a pseudo GUI 560 is shown that allows members to enter information about themselves in order to create a personal data record. Users can enter information in this GUI in various data fields. In the preferred embodiment, these fields include: Name 560-2, Home Address 560-4, Home Phone 560-6, Work Address 560-8, Work Phone 560-10, Birthday 560-12, High School 560-14, Year of High School Enrollment 560-16, High School Graduation Year 560-18, College 560-20, Year of College Enrollment 560-22, and College Graduation Year 560-24.

In certain of these data fields, the user can specify groups with which he wishes to affiliate himself, and the beginning and ending dates of the affiliation. In the preferred embodiment, the data fields High School 560-14 and College 560-20 represent categories of groups. In the data field Year of High School Enrollment 560-16, the user enters the beginning date of the affiliation with the group specified in the data field High School 560-14. In the data field High School Graduation Year 560-18, the user enters the ending date of the affiliation with the group specified in the data field High School 560-14. In the data field Year of College Enrollment 560-22, the user enters the beginning date of the affiliation with the group specified in the data field College 560-20. In the data field College Graduation Year 560-24, the user enters the ending date of the affiliation with the group specified in the data field College 560-20. In both of these cases, the beginning date and ending date establish a date range during which time the user was affiliated with the group in question.

Once the user of the client computer 370 (FIG. 5) enters information in each data field in the GUI 560 shown in FIG. 7, he clicks the Submit button 560-26 (or performs some equivalent action) and the information entered is transferred via the computer communications network 360 (FIG. 5) to the server computer 330, where the server personal contact manager software 343 stores the information in the appropriate tables 350 of a database 340.

Referring now to FIG. 8, a pseudo GUI 580 is shown that allows a first user to select other users they wish to add to their personal address book. The list of contacts is created based on the group affiliation information the first user enters in the data fields College 560-20, Year of College Enrollment 560-22, and College Year of Graduation 560-24 in the Pseudo Registration GUI 560 shown in FIG. 7. A similar GUI 580 would exist for the group specified in the data field High School 560-14 in the pseudo 560 GUI shown in FIG. 7.

In each version of the GUI 580 shown in FIG. 8, a text description 580-2 at the top of the GUI explains to the first user that other members have been found who had the same affiliation as the first user during the same period of time as the first user. The name 580-6 of the group in which the first and second users share an affiliation is displayed and the date range 580-8 of the first user's affiliation with that group is displayed.

If a second user whose personal information is stored in the tables 350 of the database 340 on the server computer 330 has specified the same group affiliation as that specified by the first user in the College 560-20 data field, and that second user has specified a date range for that affiliation that intersects with the date range specified by the first user in the Year of College Enrollment 560-22 and College Graduation Year 560-24 data fields, the Name 580-10 of the second user and the ending date 580-12 of the second user's affiliation with that group are displayed.

A second text description 580-4 at the top of the GUI 580 instructs the first user to select any of the second users listed whom the first user wishes to add to his personal address book. If the first user wishes to add a second user to his personal address book, the first user clicks the checkbox 580-14 to the left of the Name 580-10 (e.g., "John Doe") for that second user. Once the first user has finished specifying the users he wants to add to his address book, he clicks the Submit button 580-16, and the information entered is transferred via the computer communications network 360 to the server computer 330 where it is stored in the appropriate tables 350 of the database 340.

A pseudocode description of the actions performed by the personal contact manager software 343 to display the group member list is shown in Appendix A. This pseudocode fragment (and the others that follow) is written in a structured English that is similar to computer languages such as Pascal, FORTRAN and C. The pseudocode fragments are not described herein as they are self-explanatory. The tables and fields referred to in the pseudocode fragments correspond to the tables and fields described in reference to FIG. 6.

Referring now to FIG. 9, a pseudo GUI 600 is shown allowing a first user to specify which types of data fields from the first user's personal data record to grant a specific second user permission to view. If a first user specifies a second user whom the first user would like to add to his personal address book, as explained in the description of FIG. 8, the second user will receive notification (issued by the contact manager program 343—FIG. 5) that the first user has "linked" to him. If the second user chooses to return the link to the first user, the preferred embodiment of the present invention will display the pseudo GUI 600 shown in FIG. 9 with the name of the first user 600-5, allowing the second user to set data field permissions for the first user. Unlike the prior art, which does not allow the first user to specify data field permissions for individual other users, the preferred embodiment of the present invention allows the first user to specify permissions separately for each individual other user in whose personal database the first user has chosen to be included.

A text description 600-2 at the top of the pseudo GUI in FIG. 9 instructs the first user to specify which types of data fields from the first user's personal data record to allow to appear in the personal address book of the second user, whose name 600-4 is shown below. Several types of data field permission are listed, each with a check box to the left enabling the first user to select or deselect the permission type. For example, to grant the second user 600-4 permission to view the information from the first user's personal data record indicated by the permission type denoted "Crossing Paths Notification Permission," the first user would check the box 600-7 to the left of the permission type Crossing Paths Notification Permission 600-6. To deny the second user 600-4 permission to view the information from the first user's personal data record indicated by the permission type denoted "Personal Information," the first user would uncheck the box 600-9 to the left of the permission type Personal Information 600-8.

In the preferred embodiment of the present invention, the levels of permission are as follows: Crossing Paths Notification Permission 600-6, Personal Information 600-8, Work Information 600-10, Birthday Notification 600-12, and Friends of Friends Information 600-14. However, the present invention is not limited to the levels of permission shown in the preferred embodiment. The present invention is flexible to allow permission categories to be modified as needed.

Each permission type allows the second user to view information from the first user's personal data record in specific data fields, according to a specific set of rules. In the preferred embodiment of the present invention, these permission rules are as follows:

If member A links to member B, member A can grant any of the permissions discussed below to member B.

Even if member B does not reciprocate the link to member

A, an e-mail forwarding address for member B will be included in the Virtual Address Book for member A.

For example, the e-mail address

"memberB@planetall.com," which maps to the actual e-mail address that member B has entered into his/her own record, will appear in member A's Virtual Address Book, but nothing else.

When member A first links to member B, member B is notified on the Web site and in an e-mail.

If member B elects not to grant any permissions to member A, member A will not appear in member B's Virtual Address Book.

If member B grants any permissions to member A, a listing in member B's Virtual Address Book will be created for member A, and the listing will contain whatever information member A has given permission for member B to see.

If member B grants Personal Information 600-8 permission to member A, member B's home address and phone number (if available) will appear in member A's Virtual Address Book and member A will be informed when member B changes the relevant information in his/her own listing.

If member B grants Work Information 600-10 permission to member A, member B's work address and phone number (if available) will appear in member A's Virtual Address Book and member A will be informed when member B changes the relevant information in his/her own listing.

If member B grants Crossing Paths Notification Permission 600-6 to member A, member A will be able to be informed when member B will be in the same city as member A. If member A and member B are both based in the same city, member A will only be informed when member A and member B are traveling to the same destination.

If member B grants Birthday Notification 600-12 permission to member A, member B's birthday and anniversary (if available) will appear in member A's Virtual Address Book and member A will be notified when member B's birthday or anniversary are approaching.

If member B grants Friends of Friends Information 600-14 permission to member A, if member A searches for information about the contacts of his/her contacts, such as who lives in a particular city or is associated with a particular group, information from member B's circle of contacts will be included in the search results, if applicable.

Either member can modify permissions at any time. Either member can delete the other member as a contact at any time.

Pseudocode descriptions of the actions performed by the personal contact manager software 343 to display address information of contacts and to perform birthday and address change notifications are shown in Appendices B, C and D, respectively. Each of these operations depends on which permissions respective users have been granted by the owner of the information.

Once the first user has finished specifying the data field permissions for the second user 600-4, he clicks the Submit button 600-16 and the information entered is transferred via the computer communications network 360 to the server computer 330 where it is stored in the appropriate tables 350 of the database 340 (see FIG. 5). A pseudocode description of the actions performed by the personal contact manager software 343 to enable a user to change the permissions of contacts is shown in Appendix H.

Referring now to FIG. 10, a pseudo GUI 618 that displays the information stored in a user's personal address

book is shown. The information in a user's personal address book is stored in the appropriate tables 350 of the database 340 on the server computer 330, to which the client computer 370 is connected via the world wide web 360 (see FIG. 5).

5) The information in each user's personal address book is customized for that user, as described below. Each first user's personal address book contains information about each second user who has given the first user permission to view information in the second user's personal data record

10 636. Which categories of each second user's information are displayed in the first user's personal address book is controlled completely by the second user, as explained in the description of FIG. 9. In addition, each second user's information is entered and maintained completely by the

15 second user (e.g., "Donald Tully"), as explained in the description of FIG. 7.

FIG. 10 illustrates the Address Book pseudo GUI 620 at three levels, in which each next level allows the user to view progressively more detail about the contacts in his personal address book.

In level 1 620 of the pseudo address book GUI 618, each letter of the alphabet is shown 622. By clicking on any letter of the alphabet 622, a first user can display a listing of the contacts whose last names begin with the letter of the

25 alphabet selected, and about whom information is stored in the first user's personal address book. This information is displayed in the level 2 626 of the pseudo address book GUI.

For example, if the first user clicks on the letter "T" 624 in level 1 620 of the pseudo address book interface, all contacts 30 whose last names begin with the letter T and about whom information is stored in the first user's personal address book will be displayed 628 in level 2 626 of the pseudo address book GUI.

In level 2 626 of the pseudo address book GUI, a listing 35 of the second users whose last names begin with the letter of the alphabet selected in level 1 620 of the pseudo address book GUI, and about whom information is stored in the first user's personal address book, is shown. By clicking on any second user's name, the first user can display the information

40 about that second user stored in the first user's personal address book. This information is displayed in level 3 632 of the pseudo personal address book GUI. For example, if the first user clicks on the second user name "Tully, Donald" 630, the information pertaining to Donald Tully stored in the first user's personal address book will be displayed in level 3 632 of the pseudo address book GUI.

In level 3 632 of the pseudo address book GUI, information 35 634 is shown about a specific second user that is stored in a first user's personal address book. Only the categories 50 of information from the second user's personal data record that the second user gave the first user permission to view are displayed. The second user's information is entered and maintained completely by the second user.

In the preferred embodiment of the present invention, the 55 second user's e-mail address 634-2 is displayed if the second user gave the first user any type of data field permission. The second user's work address and phone number 634-4 are displayed only if the second user gave the first user Work Information permission. The second user's home address 60 and phone number 634-6 are displayed only if the second user gave the first user Personal Information permission. The second user's birthday and birth year 634-8 are displayed only if the second user gave the first user Birthday Notification permission. These permission rules are simply 65 examples from the preferred embodiment. The present invention is not limited to the permission rules used in the preferred embodiment.

A pseudocode description of the actions performed by the personal contact manager software 343 to display the address book listing is shown in Appendix B.

Referring now to FIG. 11, a member update pseudo GUI 650 is shown. This pseudo GUI 650 provides a first user with specific information that has changed about the other users to which the first user is linked, plus new information about contacts to whom the first user may wish to link. The information displayed in a user's member update is stored in the appropriate tables 350 of the database 340 on the server computer 330, to which the client computer 370 is connected via the world wide web 360. The member update pseudo GUI 650 is automatically displayed on the user interface 380 of the user workstation 370, at an interval preset by the user. For example, FIG. 11 displays a hypothetical member update 650-2 released on Dec. 7, 1998. The information displayed in the data fields below is information that has changed between Dec. 7, 1998 and the date of the previous update, the interval between which has been previously specified by the user. The information shown in each user's member update is customized for that user, as described below.

In a first portion of the member update pseudo GUI 650 shown in FIG. 11, if one or more of the second users who have linked to a first user and have provided Birthday Notification permission to the first user have upcoming birthdays, a text description 650-4 alerts the first user of the upcoming birthday(s). The names and birthdays 650-6 for those second users are listed below. In the preferred embodiment of the present invention, the first user will receive this notification 2 weeks, 1 week, 2 days, and 1 day in advance of a particular upcoming birthday, and on the actual date of the birthday. The first user does not need to collect and input the birthday dates for each second user who has linked to the first user. Each second user's birthday information is entered and maintained completely by the second user, as shown in the Birthday field 560-12 of FIG. 7, the registration form pseudo GUI 560.

In another portion of the member update pseudo GUI shown in FIG. 11, if one or more of the second users who have linked to the first user and have provided Personal Information permission or Work Information permission to the first user have changed their work or home address, a text description 650-8 alerts the first user. If a second user has changed his work address information and has given the first user Work Information permission, the second user's new work address information 650-10, 650-12 is displayed. If a second user has changed his home address information and has given the first user Personal Information permission, the second user's new home address information is displayed.

Each second user's address information is entered and maintained completely by the second user, as shown in the registration form pseudo GUI 560 of FIG. 7. After changing his address information in his personal data record, the second user does not need to specify that the new address information be provided to each first user to whom he has linked and given the proper form of data field permission. The new address information is provided to each first user quickly and automatically. In addition, the architecture of the present invention is scalable to include millions of users.

In another portion of the member update pseudo GUI 650 shown in FIG. 11, if one or more members has affiliated with a group with which the first user is also affiliated, a text description 650-14 will alert the first user. The name of the second user, the name of the group in which the first and second users share an affiliation, and the ending date of the second user's affiliation with that group are displayed 650-16.

This portion of the registration form pseudo GUI 650 functions similarly to the group list form pseudo GUI shown in FIG. 8. If a new second user who fills out a registration form such as the pseudo GUI in FIG. 7, and therefore whose personal information is stored in the tables 350 of the database 340 on the server computer 330 has specified the same group affiliation as that specified by the first user in the College 560-20 data field, and that second user has specified a date range for that affiliation that intersects with the date range specified by the first user in the Year of College Enrollment 560-22 and College Graduation Year 560-24 data fields, the Name of the second user and the ending date of the second user's affiliation with that group 650-16 are displayed. Similarly, if the first user and the new second user were affiliated during an intersecting period of time with the group specified in the data field High School 560-14 in the pseudo GUI 560 shown in FIG. 7, the Name of the second user and the ending date of the second user's affiliation with that group 650-16 are displayed.

20 A pseudocode description of the actions performed by the personal contact manager software 343 to display a list of service members who have recently joined a user's groups (i.e., members who are not current contacts of the user) is shown in Appendix E.

If the first user wishes to add contact information to his personal address book for any of the second users listed 650-16, the first user can do so in a GUI similar to the group list form pseudo GUI 580 shown in FIG. 8. Each second user to whom the first user has initiated a link will then be informed of the link, and can then return the link and specify data field permissions for the first user, if any, as explained in the description of FIG. 9.

25 A pseudocode description of the actions performed by the personal contact manager software 343 to identify people 35 who have linked to a particular user are shown in Appendix F.

In another portion of the member update pseudo GUI 650 shown in FIG. 11, if a second user has initiated a link to a first user, the first user will be automatically notified 650-18 40 that a link has been made. For each second user that has initiated a link, the user's name 650-20 is shown. If the first user wishes, the first user can then return the link and specify data field permissions for the second user, if any, as explained in the description of FIG. 9.

Another section 650-22 of the member update pseudo GUI 650 shown in FIG. 11 is used to inform a first user when the travel plans he has entered into the system overlap with the travel plans that any of his contacts has entered into the system, as long as the contact has granted the first user Crossing Paths Notification permission.

This system, termed "Crossing Paths Notification" in the preferred embodiment of the present invention, operates as follows. The home city or "base city" for each user is determined from information entered by that user in the Home Address data field 560-4, as explained in the description of FIG. 7. The "City" table 550 (FIG. 6) stored on the server computer 330 includes 1.7 million names of cities around the world. Each of these cities is associated with a precise latitude and longitude. If the user's base city cannot be matched to a city in the "City" table, the user can add the new city to the "City" table by giving the name of another city that is already in the "City" table that is nearby the user's base city. The user's base city is assigned the same latitude and longitude as the existing city. This information is used to associate each user with a precise longitude and latitude, and determine all cities within a 29-mile radius of the user's base city.

Whenever a user is planning to travel, he can specify the dates during which he will be away and the city he will be visiting. If a second user has granted a first user Crossing Paths Notification permission, and the first user has entered a Travel Event to a city that is within a 29-mile radius of the base city of the second user, the first user will be notified 650-22 (FIG. 11) that he will be crossing paths with the second user 650-24 (e.g., "Andrew Kress"), as long as the second user has not also scheduled a travel event for the same time period. In another scenario, if a second user has granted a first user Crossing Paths Notification permission, and the first user has entered a travel event to a city that is within a 25-mile radius of a city to which the second user has scheduled a travel event during the same time period, the first user will be notified 650-22 that he will be crossing paths with the second user 650-24. Travel events are described more fully in reference to FIG. 12.

The Crossing Paths Notification system is able to handle multiple cities in a single day. For instance, if a first user lives in Boston but is traveling to New York on March 5, then the first user will be informed if any contacts will be crossing paths on that day in either city. In addition, this system is scalable to millions of users. A pseudocode description of the actions performed by the personal contact manager software 343 to enable a user to receive crossing paths notification is shown in Appendix I.

The final section 650-26 of the member update pseudo GUI 650 shown in FIG. 11 is used to inform a first user which of his contacts has an astrological sign compatible with that of the first user on the date of the member update. Each member is associated with one of the twelve astrological signs based on the information he entered in the Birthday data field 560-12 in the registration form pseudo GUI 560 shown in FIG. 7. Each day of the year is mapped to one of these twelve signs. This information is stored in the appropriate table 350 in the database 340 on the server computer 330. On a given day, all of a member's contacts who are associated with "sign of the day" are deemed to be compatible with the member. Only the names of contacts who have given the first user Birthday Notification permission will be shown in the member update pseudo GUI for the first user. A pseudocode description of the actions performed by the personal contact manager software 343 to enable a user to receive notification of compatible contacts is shown in Appendix J.

The permission rules used in reference to FIG. 11 are simply examples from the preferred embodiment. The present invention is not limited to the permission rules used in the preferred embodiment.

Referring now to FIG. 12, a pseudo Add Travel Form GUI 660 and a pseudo Crossing Paths List GUI 670 are shown. These two screens are used in the Crossing Paths Notification System. If a first member is planning a trip, the first member can use the pseudo Add Travel Form GUI 660 to add a Travel Event, in which he specifies the location 660-2, 660-4, 660-6, dates 660-8, 660-10, and contact information 660-20 for the intended trip. In the pseudo Crossing Paths List GUI 670, the first member is informed which of the second members to whom he is linked and who have granted him Crossing Paths Permission will be in the vicinity of the city to which the first user is travelling, during the time period of the specified Travel Event. The first user can then use the pseudo Crossing Paths List GUI 670 to select which of the displayed second users the first user would like to inform of the first user's specified Travel Event.

The pseudo Add Travel Form 660 is displayed on the user interface 380 (FIG. 5) of a user's client computer 370 when

the user chooses to schedule a Travel Event. The user enters information about his scheduled trip in the data fields shown. In the Traveling To City field 660-2, the user enters the name of the city to which he is traveling. In the State field 660-4, the user enters the name of the state in which is located the city to which he is traveling. In the Country 660-6 field, the user enters the name of the country in which the city to which he is traveling is located. The information entered in these three fields 660-2, 660-4, 660-6 is used to locate the city for the Travel Event in the City table 550 on the server computer 330. The exact latitude and longitude of the Travel Event city is then determined and a list is created of all cities located within a 25-mile radius of the Travel Event city. In the Arrive in City on Date field 660-8, the user enters the first date on which he will be in the Travel Event city. In the Leave City on Date field 660-10, the user enters the date beginning on which he will no longer be in the Travel Event city. The information entered in these two fields 660-8, 660-10 is used to determine the date range for the Travel Event. Finally, in the How to Get in Touch While in This City data field 660-20, the user enters the method for contacting him during the Travel Event. After the user has finished entering information in the pseudo Add Travel Form GUI (12-1), the information entered is stored by the personal contact manager 343 in the Travel\_Event table 520 on the server computer 330.

The pseudo Crossing Paths List 670 is displayed on the user interface 380 of the first user's client computer 370 after a first user has scheduled a Travel Event user the pseudo Add Travel Form 660. A text message 670-2 issued by the personal contact manager 330 informs the first user that one or more of his contacts will be in the same city as the first user during the first user's scheduled Travel Event. Those contacts (e.g., Scott Ulem, Taylor Pierce, Betsy Klein) who live in the city of the first user's scheduled Travel Event are listed 670-6, as well as those contacts (e.g., Tania Gutsche) who will be visiting the city of the first user's scheduled Travel Event 670-8. The contacts listed in the field 670-6 are those second users who have granted the first user Crossing Paths Permission, and who have listed in the Home Address field 560-4 (FIG. 7) of their Personal Data Record the city of the first user's scheduled Travel Event, or any city within a 25-mile radius of the first user's scheduled Travel Event. The contacts listed in the field 670-8 are those second users who have granted the first user Crossing Paths Permission, and who have scheduled a Travel Event to the city of the first user's scheduled Travel Event, or any city within a 25-mile radius of the first user's scheduled Travel Event, during the date range of the first user's scheduled Travel Event. For each contact name listed in both fields 670-6, 670-8, the first user can choose to inform that contact of the first user's scheduled Travel Event by clicking on the checkbox to the left of that contacts name. When the first user is finished selecting contacts, he then clicks the Submit button 670-10, which copies the information entered to the server computer (5-45) to be stored in the tables 350 by the networked personal contact manager 343. For each second user whom the first user selected, the second user is informed, in a screen similar to the pseudo Member Update GUI shown in FIG. 11, of the first user's Travel Event and the means of contacting the first user 660-20 during the Travel Event.

Referring now to FIG. 13, a diagram illustrating the Friends of Friends system is shown. The Friends of Friends system allows a first member to search for the names of contacts of their contacts who live in the same city as the first member are affiliated with a group with which the first member is also affiliated. When a user performs a Friends of

Friends search, the personal contact manager 343 displays, via the web server software 342, the results of the search on the user interface 380 (FIG. 5) of the first user's client computer 370 in a GUI similar to a pseudo Friends of Friends report GUI 688. After locating a second member with a friend of a friend, the first member can then link to that second member in order to add the second member to the first user's Personal Address Book, as explained in the descriptions of FIG. 8 and FIG. 9 above.

In the preferred embodiment of the personal contact manager 343, the Friends of Friends system operates as follows. If a Member A 680 is linked to a Member B 682 with any level of permissions 681 and the Member B 682 is linked to a Member C 684 with any level of permissions 685, then if Member C 684 grants to Member B 682 Friends of Friends permissions 687 and Member B 682 also grants to Member A 680 Friends of Friends permissions 683, then Member A is eligible to receive Friends of Friends notification about Member C. When a first user performs a Friends of Friends search, the results of the search will include all second users who have affiliated themselves with a group with which the first user is affiliated and all second users who live in the same city in which the first user lives, so long as the first user is eligible to receive Friends of Friends notification about those second users, as described above. For example, if Member A and Member C both belong to Group A 686, and Member A is eligible to receive Friends of Friends notification about Member C, then the result of Member A's Friends of Friends search 688 generated by the personal contact manager 343 will include Member C 690.

A pseudocode description of the actions performed by the personal contact manager software 343 to perform a search for friends and friends of friends in a specific city is shown in Appendix G.

The present invention is not limited to the search criteria or levels of separation in the preferred embodiment. The database architecture in the present invention is flexible to allow searches to be extended to more than one degree of separation. For instance, it would be possible to add a Friends of Friends of Friends search feature. The architecture is also flexible to allow new search criteria to be added.

Referring to FIG. 5, in each of the embodiments described above, the user information is stored on the server 330 and all user access to the user information mediated by a client web browser 382, the web server software 342 and the server personal contact manager software 343. In an alternative embodiment, which is configured for personal information managers (PIMs), such as the U.S. Robotics Palm Pilot, a user is able to synchronize their user information and their PIM database 390 through an importation/synchronization function performed by the personal contact manager software 343. The synchronization operation can be performed in either direction (i.e., client to server or server to client). The server personal contact manager software 343 will then use the web server software 342 to communicate with the PIM software 392 of the user's contacts, if applicable, and, in accordance with the permission scheme already described, synchronize the databases 390 in the contacts' PIMs. All database, personal contact management and linking operations already described are operable in the alternative embodiment, except the GUIs might be different, depending on the graphical capabilities of the client 370 running the PIM program 392. Thus, the alternative embodiment allows full synchronization of PIMs and the server database 340.

A data flow diagram illustrating the operation of the alternative embodiment is shown in FIG. 14. In the illus-

trated situation a user A submits an address change from their client computer 370A. In response to the update, the personal contact manager 343 running on the server 330 updates user A's address information in the server database 340 (not shown) and issues an update notification to the client computer 370B used by user B, who is a contact of user A. This alternative embodiment assumes that user B has a PIM (also referred to as a personal digital assistant or PDA) that they would like to synchronize with the server database 340. In such a case PIM Software 392 running on the client 370B performs the synchronization operation based on the user A address update information provided by the server 330. Following the synchronization operation, the PDA database 390 has the same information for user A as the server database 340. Alternatively, the PDA 750 can be coupled directly to the Internet (indicated by the dashed line), in which case it operates substantially as a typical client computer 370 described in reference to FIG. 5. However, one difference is that the PDA 750 maintains its own database 390 instead of relying solely on the server database 340.

While the present invention has been described with reference to a few specific embodiments, the description is illustrative of the invention and is not to be construed as limiting the invention. Various modifications may occur to those skilled in the art without departing from the true spirit and scope of the invention as defined by the appended claims.

## APPENDIX

### Display Group Member List

```

Submit group name.
Match group name to GroupID in Group table.
Join Affinity table to Customer table and CustomerPrefs
table based on CustomerID.
Show contact information from Customer and Customer-
Prefs tables when the the Affinity table contains a record
matching the CustomerID to the specified GroupID.

```

## APPENDIX B

### Display Address Book Listing

```

if MemberFriend AND a Reciprocated Link then
If have Personal or Professional Permissions then
  Show person's real email address.
else
  Show person's PlanetAll address.
end if
If have Personal Permission and Biography Exists then
  Show Biography.
end if
If (have Personal or Professional Permissions) AND uni-
  versal resource locator (URL) exists then
  Show URL.
end if
If (Person is visible in group and Group Perms>0) AND
  you have Common Groups then
  Show the groups you have in common.
end if
If Person is in one of more of your personal email lists
  then
  Show the lists the person belongs to.
end if

```

**19**

First Name from the Customers table.  
 Last Name from the Customers table.  
 Record Date from the Friends table.  
 Permission level from the Friends table.  
 Where my CustomerID is not among the CustomerIDs found in the following search:  
 CustomerID in the Friends Table is my CustomerID  
 AND the Record Date from the Friends table is within the last 30 days  
 AND I haven't already linked to the person

**APPENDIX G**

Search for Friends of Friends in a Particular City

Specify City. Match to CityID in City table.  
 Create a list of all my contacts  
 Go to the Friend table and select all Customers where FriendID=my CustomerID  
 Make a temporary table linking the Friends table to itself called Friend\_1 and establish the following relationships:  
 CustomerIDs for the contacts of my contacts appear in the Customer field of the Friends table  
 CustomerIDs for my contacts appear in the Friend field of the Friends table  
 CustomerIDs for my contacts also appear in the Customer field of the Friends\_1 table (this is how the tables are joined)  
 My Customer ID appears in the Friend field of the Friend\_1 table  
 The Friend and Friend\_1 tables are joined on t  
 Show information for the contacts of my contacts (i.e. the Customers from the Friends table) where the following conditions are true:  
 The Friends of Friends permission was granted from the contacts of my contacts to my contacts.  
 The Friends of Friends permission was granted from my contacts to me.  
 The contact does not already appear in the list of all my contact created above.  
 The city for the contact of my contact matches the specified city.

**APPENDIX H****Change Permissions**

Join the Customer table to the Friend table based on CustomerID.  
 Create a list of all my contacts:  
 Go to the Friend table and select all Customers where FriendID=my CustomerID.  
 Show First Name and Last Name for my contacts from the Customer table.  
 Allow me to pick a name from this list as the contact whose permissions I would like to change.  
 Display the permission level that I have given this contact.  
 It is stored as the PermissionType field in the Friend table.  
 Allow access to the PermissionType for this record in the Friend table.

**APPENDIX I****Crossing Paths Notification**

Create a list of all my contacts:  
 Go to the Friend table and select all Customers where FriendID=my CustomerID.

**20**

Create a list of all my contacts' travel events:  
 Go to the Queue Travel Event table and select all Queuels where the CustomerID is in my list of contacts.  
 Do not include trips for people who linked to me but did not give me crossing paths permission:  
 Check the permissions field in the Friend table for each of my contacts to see if I have crossing paths permissions.  
 Do not include trips if my contact specified that I should not be informed:  
 For each of my contacts' trips, check the Travel Exception table to see if my CustomerID is included in the list of people who should not be informed of the trip.

Create a list of my location for the next seven days:  
 Select the arrival date, departure date, and city for all my trips in Queue Travel Event for the next seven days.  
 For days when I am not travelling, select my city from the customers table.  
 Select from the list of my friends' trips, all the trips to cities that are within 3000 latitude and 3000 longitude to my location for each of the next seven days.  
 Go to the Customers table and find the names of all the people with whom I will be crossing paths.  
 Select first name and last name from the customers table for all the CustomerIDs in the list of my contacts trips

**APPENDIX J****Compatible Contacts**

Create a list of all my contacts:  
 Go to the Friend table and select all Customers where FriendID=my CustomerID.  
 Determine my Zodiac sign:  
 Select my DayOfYear from the Customers table.  
 Select the Zodiac sign from the Zodiac table where my DayOfYear is between the DayFrom and DayTo fields.  
 Determine my compatible Zodiac sign for today:  
 Go to the Horoscope table and select the Compatible field from the row for my Zodiac sign and today's date.

Find my compatible contacts for today:  
 Select the DayFrom and DayTo fields from the Zodiac table for my compatible zodiac sign.  
 Select my contacts from the list of all my contacts whose DayOfYear is between the DayFrom and DayTo fields for my compatible sign.

What is claimed is:  
 1. A networked contact management system, comprising:  
 a database which contains personal records of multiple users, each record including multiple fields, including fields for storing personal contact information; and  
 contact manager software which provides restricted access to the database through an interface in which (1) users select other users from the database to include in their own, respective virtual personal address books without the need to enter information for such users, (2) if a first user selects a second user to include in the first user's virtual personal address book, the second user is provided an option to specify the types of information

of the second user's personal record to be viewable by the first user, (3) users directly update their own respective personal records within the database, and (4) the personal records stored within the database are at least partially viewable as virtual address book entries, so that updates made by users to their own respective personal records are reflected automatically within the virtual personal address books of other users without the need to propagate or separately apply the updates to individual address books.

2. The networked contact management system as in claim 1, wherein the interface further provides the second user an option to grant permission to the first user to be notified of at least one type of occasion.

3. The networked contact management system as in claim 2, wherein the at least one type of occasion includes at least one of the following: (a) a birthday of the second user, (b) an anniversary of the second user, and (c) a correspondence in schedules of the first and second users.

4. The networked contact management system as in claim 1, wherein the database contains information about a plurality of affinity groups, and the contact manager software implements functions for (a) allowing the multiple users to select and join individual affinity groups, (b) notifying a new member of an affinity group of existing members of the affinity group, and (c) presenting the new member an option to select individual members of the affinity group to add to the new member's personal address book.

5. The networked contact management system as in claim 4, wherein the contact manager software automatically notifies existing members of the affinity group of new members.

6. The networked contact management system as in claim 1, wherein the contact manager software receives and stores information about travel plans of individual users, and uses the information about travel plans to detect that a user and a contact of the user will be in the same location at the same time.

7. The networked contact management system as in claim 1, wherein the contact manager software automatically notifies the first user when the second user modifies a datum of the second user's personal record that the first user is permitted to view.

8. The networked contact management system as in claim 1, further comprising synchronization software which synchronizes a personal digital assistant (PDA) device with a user's virtual personal address book over a computer network.

9. A networked personal contact management system, comprising:

- a database which contains personal data records of a plurality of users, at least some of the data records including contact information of respective users; and
- a server system which provides restricted access to the database through an interface that provides functions

for each user to at least (a) directly modify the user's own repetitive personal data record within the database, (b) select other users from the database to add to a virtual personal address book of the user, and (c) specify, on a user-by-user basis, permissions for other users to view the personal data record of the user through virtual personal address books of such other users;

wherein users directly view the data records of other users through the virtual address books according to said permissions, so that updates by users to their own respective personal records are reflected automatically within the virtual personal address books of other users.

10 10. The networked personal contact management system as in claim 9, wherein the interface allows each user to grant said permissions separately for each of a plurality of data types.

15 11. The networked personal contact management system as in claim 9, wherein the server system notifies the user of at least one of the following types of occasions associated with a contact of the user: (a) a birthday of the contact, (b) an anniversary of the contact, and (c) a correspondence in schedules of the user and the contact.

20 12. The networked personal contact management system as in claim 9, wherein the database contains information about a plurality of affinity groups, and the server system implements functions for (a) allowing the plurality of users to select and join individual affinity groups, (b) notifying a new member of an affinity group of existing members of the affinity group, and (c) presenting the new member an option to select individual members of the affinity group to add to the new member's virtual personal address book.

25 13. The networked personal contact management system as in claim 12, wherein the server system automatically notifies existing members of the affinity group of new members.

30 14. The networked personal contact management system as in claim 9, wherein the server system receives and stores information about travel plans of individual users, and uses the travel plans to detect that the user and a contact of the user will cross paths.

35 15. The networked personal contact management system as in claim 9, wherein the server system automatically notifies the user when contacts recorded within the user's virtual personal address book modify their respective contact information displayed within the virtual personal address book.

40 16. The networked personal contact management system as in claim 9, further comprising synchronization software which synchronizes a personal digital assistant (PDA) device with a user's virtual personal address book over a computer network.

\* \* \* \* \*